



THE UNIVERSITY OF
WAIKATO
Te Whare Wānanga o Waikato

Research Commons

<https://researchcommons.waikato.ac.nz/>

Research Commons at the University of Waikato

Copyright Statement:

The digital copy of this thesis is protected by the Copyright Act 1994 (New Zealand).

The thesis may be consulted by you, provided you comply with the provisions of the Act and the following conditions of use:

- Any use you make of these documents or images must be for research or private study purposes only, and you may not make them available to any other person.
- Authors control the copyright of their thesis. You will recognise the author's right to be identified as the author of the thesis, and due acknowledgement will be made to the author where appropriate.
- You will obtain the author's permission before publishing any material from the thesis.

Process Integration and Electrification with Digital Twins

A thesis

submitted in fulfilment

of the requirements for the degree

of

Doctorate of Philosophy in Engineering

At

The University of Waikato

by

Benjamin James Lincoln



2026

Abstract

The decarbonisation of industrial process heat is one of the most pressing challenges in the global energy transition. In New Zealand, fossil fuels remain the dominant source of process heat, despite having over 80% renewable electricity generation. Milk powder production is a major consumer of process heat, with evaporation and drying processes relying on large amounts of coal- and gas-fired steam. Electrification technologies such as industrial heat pumps and mechanical vapour recompression (MVR) have the potential to significantly reduce emissions, yet widespread adoption has been limited because of the complex interactions between heat and power, in addition to uncertainties around practicality. Conventional process integration (PI) techniques were designed for fossil-fuelled utilities and are poorly aligned with the work requirements and integration constraints of electrification. Meanwhile, legacy simulation tools are ill-suited to the complex fluids and system interactions of food and dairy processes.

This thesis addresses these gaps by developing a generalisable Process Integration and Electrification (PI&E) methodology that combines exergy-based targeting, retrofit strategies, and techno-economic evaluation coupled with an iterative design-centric digital twin framework. The thesis is structured in two parts. Part A develops the digitalisation foundations, including the preparation of a milk evaporation case study, the creation of advanced thermophysical property packages for complex fluids (milk, refrigerants, humid air), and the construction of a design digital twin using both commercial and open-source platforms. Part B applies the digital twin to PI&E, integrating operational optimisation, Exergy Pinch Analysis, and systematic evaluation of electrification technologies in both greenfield and retrofit contexts.

For greenfield design, the research extends Pinch Analysis principles to heat pump integration by utilising heat pockets to create multiple Pinch points, enabling systematic minimisation of temperature lift and improved integration opportunities. Building on this, an iterative PI&E design workflow was developed to guide technology placement and evaluate electrification pathways. This culminated in the design of a novel fully electric milk evaporator system that achieved a specific electricity consumption of 120 kWh per tonne of milk powder, compared with 159 kWh/tp for a simpler single heat pump design, demonstrating higher efficiency.

For retrofit applications, the thesis advances PI&E by extending heat pump bridge analysis to explicitly include process unit heat flows, allowing process modifications to be considered alongside heat exchanger reconfiguration. This innovation addresses a key gap identified in previous literature,

enabling more retrofit strategies. The method was demonstrated through multiple related case studies of milk evaporator plants, producing a set of common retrofit solutions. These include replacing thermal vapour recompression (TVR) and/or direct steam injection with MVR systems, which were shown to deliver lower levelised costs of heat compared with reference boiler-based designs.

The culmination of the research is a unified PI&E methodology that combines digital twins, rigorous thermodynamic analysis, and practical integration strategies. The results show that electrification of milk evaporation systems can be achieved in both new and existing plants with significant efficiency gains and competitive economics. PI&E has been tested across multiple platforms: Aspen HYSYS, DWSIM and the Ahuora Digital Twin Platform, powered by IDAES – proving to be a platform-agnostic, yet digitalisation-centred, methodology. Although developed and applied in the context of New Zealand’s dairy sector, the methods and insights are broadly transferable to other low- to medium-temperature process industries, offering a robust and scalable pathway to accelerate industrial decarbonisation.

Acknowledgements

First and foremost, I thank God for the strength, wisdom, and blessings that sustained me throughout this journey. There were many moments of doubt and difficulty, yet I found peace and courage in His presence. *“I can do all things through Christ who strengthens me”* (Philippians 4:13) has been a guiding truth during this work, and I am grateful beyond words.

To my Mum and Dad, Lissete and Wayne Lincoln, thank you for every sacrifice and every moment of encouragement that brought me to this point. Your belief in me has been unwavering, and I owe this achievement to the foundation you built. To my sister, who has been both a lifelong friend and a source of support, from toddler adventures to flatmate days, thank you for always being there.

I am indebted to my supervisors, Martin, Michael, and Brent, for their guidance, coaching, and support in developing both my technical and professional skills. Your advice, patience, and encouragement gave me confidence, and I knew you always had my back.

To my chief supervisor, Tim Walmsley: thank you for the incredible number of opportunities you created for me and for throwing me into the deep end more times than I can count. You always believed in me, even at times when I doubted myself, and that belief has shaped the researcher and person I have become.

The Ahuora office has been far more than just a workplace. Every member has had an impact on this journey. From the students I taught to those who ended up teaching me, you were all willing to engage with my unusual hypotheticals and made each day full of humour and camaraderie. A special thanks to the original crew, Elsa, Matthew, Alyssa, Keegan, Sara, and Jack, who were the springboard for this adventure. And to Stephen, my Software Platform Co-Leader: here’s to many more hours of frustrated staring at the same screen together.

Finally, the greatest thanks goes to my fiancée (and hopefully wife by the time this is read), Lana. You have been my constant support, celebrating the highs and lifting me through the lows. You always made me feel like the smartest person in the room, even when I wasn’t sure myself. They say you should never work with friends or family perhaps that advice I guess extends to never doing a PhD at the same time either, but I cannot imagine having done it any other way than with you by my side.

Contributing Publications

a. Papers directly contributing to the thesis

Several peer-reviewed publications have directly contributed to the development of this thesis. These papers collectively establish the foundations for the Process Integration and Electrification (PI&E) methodology, demonstrate their application to dairy process systems, and validate the integration of digital twin technology for both greenfield and retrofit contexts. The two central papers represent the most significant contributions, with approximately 85% of the work undertaken by the Thesis Author. Evidence of these contributions is provided in Appendix Five (Section 11.5).

1. “Process Integration and Electrification for Efficient Milk Evaporation Systems” (Lincoln et al., 2022).
2. “Process Integration and Electrification for Retrofit: Case studies of Milk Evaporator Systems” (Lincoln et al., 2024).

An additional extension of the Thesis work was published on the topic of Exergy Pinch Analysis, which builds on the exergy-based targeting methods presented in this thesis. The paper provided expanded the results for the same dairy case study, offering deeper insights into system-level work losses and integration opportunities. This publication was co-authored with the chief supervisor.

3. “Advancing Industrial Process Electrification and Heat Pump Integration with New Exergy Pinch Analysis Targeting Techniques” (T. Walmsley et al., 2024)

b. Other co-authored papers relevant to digitalisation and PI&E

4. “Process Integration And Electrification Through Multiple Heat Pumps Using A Lorenz Efficiency Approach” (Padullés et al., 2024).
5. “Retrofit Of A Galvanisation Plant Using Advanced Heat Pump Bridge Analysis” (Schlosser, Walmsley, et al., 2023).
6. “New Directions and Software Tools Within the Process Systems Engineering Ecosystem” (Burroughs et al., 2025).

Contents

Abstract	ii
Acknowledgements	iv
Contributing Publications	v
a. Papers directly contributing to the thesis.....	v
b. Other co-authored papers relevant to digitalisation and PI&E.....	v
Contents	vi
Table of Figures	xiv
List of Tables	xix
Nomenclature	xx
Abbreviations	xx
Symbols.....	xxii
Subscripts	xxiii
Units	xxiv
Prefixes.....	xxiv
Chapter 1 Introduction.....	1
1.1 Background.....	1
1.2 Thesis Aim.....	3
1.3 Structure of the Thesis	5
Chapter 2 Literature Review.....	7
2.1 Introduction.....	7
2.2 Process Heat Electrification and Industrial Decarbonisation	8
2.2.1 Process electrification case studies	10
2.2.2 Case study in context: milk powder.....	15
2.2.3 Standardised commercially available heat pump review	16
2.2.4 Commercially available mass manufactured heat pumps internationally.....	17
2.3 Process Integration for Energy Efficiency.....	17

2.3.1	Pinch analysis.....	18
2.3.2	Idealised Pinch Analysis heat pump integration methods	21
2.3.3	Current heat pump integration methods.....	22
2.3.4	The exergy per total annual cost method (EPC)	25
2.3.5	Simulation based methods	25
2.3.6	Integration for non-minimal Pinch target systems	26
2.4	Heat Pump Retrofit	26
2.4.1	Heat Exchanger Network (HEN) retrofit methods	26
2.4.2	Modified Heat Pump Bridge Analysis	28
2.4.3	Automated Heat and Work Retrofit Methods	30
2.5	Digitalisation and Digital Twins.....	31
2.5.1	A Standardised Energy Digital Twin Framework for Process Systems	34
2.5.2	Digital Twins for Milk Evaporation Plants.....	37
2.5.3	Review of Model-Based Optimization of Energy Consumption in Milk Evaporators	38
2.6	Property Modelling for Complex Fluids.....	40
2.6.1	Simulation of Milk for Evaporation.....	42
2.6.2	History of accurate property packages for vapour compression cycles.....	43
2.6.3	Current tools and challenges for implementation	45
2.6.4	Implementation of Helmholtz into the IDAES equation orientated framework	46
2.6.5	State of the art modelling of Humid Air	47
2.6.6	Formulation of the Humid Air Equation of State	48
2.6.7	Machine learning for surrogate properties packages.....	51
2.7	Conclusions.....	54
2.7.1	Gaps in the literature	55
Chapter 3	Overarching Methodology	57
3.1	Introduction.....	57

3.2	Methodology.....	58
3.3	Part A: Digital Twin Development.....	59
3.3.1	Purpose.....	59
3.3.2	Approach.....	59
3.3.3	Conclusion	60
3.4	Part B: Process Integration and Electrification Method	61
3.4.1	Purpose.....	61
3.4.2	Approach.....	62
3.4.3	Subtleties of retrofit	63
3.4.4	Conclusion	65
3.5	Conclusions.....	65
Chapter 4	Digital Twin Goals and Preparation.....	67
4.1	Introduction.....	67
4.2	Method.....	68
4.2.1	Preparing to build a digital twin.....	68
4.2.2	Defining design requirements and modelling boundaries.....	68
4.2.3	Operational requirements and design ideals	69
4.2.4	Data collection	70
4.2.5	Data for sensitivity analysis	70
4.2.6	Selecting the appropriate simulation platform	71
4.3	Case study and data analysis.....	72
4.3.1	Design intent	72
4.3.2	The milk evaporation process and design scope.....	73
4.3.3	Operational requirements and design outcomes	74
4.3.4	Data collection for the base case (Greenfield).....	75
4.3.5	Data collection for sensitivity analysis (Greenfield).....	76
4.3.6	Selection of the simulation environment	77

4.4	Conclusions.....	78
Chapter 5	Advanced Property Packages for Complex Fluids	80
5.1	Introduction.....	80
5.1.1	Equations of State: Illustration of complexity	81
5.2	The thermodynamic properties of milk	82
5.2.1	Introduction.....	82
5.2.2	Method for activity coefficient modelling for conventional simulation platforms 83	
5.2.3	Method adaptation for equation orientated frameworks.....	86
5.2.4	Results for the milk activity coefficient models	88
5.2.5	Discussion	92
5.3	Custom Helmholtz-Based Formulations for Key Pure Components.....	92
5.3.1	Introduction.....	92
5.3.2	Method for creating a Helmholtz Free Energy Property Package for Equation Oriented Programming.....	93
5.3.3	Results.....	104
5.3.4	Discussion	109
5.4	Machine-Learned Surrogate Models for Complex Mixtures.....	110
5.4.1	Introduction.....	110
5.4.2	The case study: Humid Air	111
5.4.3	Method For Surrogate Modelling of Thermophysical Properties.....	112
5.4.4	Results for simple surrogate model.....	116
5.4.5	Method for increasing surrogate performance metrics via sampling.....	119
5.4.6	Results for adaptive sampling surrogate model creation	123
5.4.7	Creating the full property package.....	127
5.4.8	Discussion	128
5.5	Conclusion	129
Chapter 6	Digital twin	132

6.1	Introduction.....	132
6.2	Methods for Creating and Evaluating Unit Models.....	133
6.2.1	Ejector / TVR modelling method.....	133
6.2.2	Direct Steam Injector (DSI)	140
6.2.3	Direct Contact Heater (DCH)	143
6.2.4	Mechanical Vapour Recompressor (MVR).....	144
6.2.5	Falling Film Evaporator	145
6.2.6	Heat Pump modelling.....	147
6.3	Methods for Deciding the Wider Solving Structures for the Total Flowsheet.....	154
6.3.1	Sequential modular simulation.....	154
6.3.2	Methods for extracting useful data for the PI&E method.....	158
6.4	Results and Discussion	160
6.4.1	Comparison of digital twin results	160
6.4.2	Gathering data	165
6.4.3	Knowledge barriers	168
6.5	Conclusion	168
Chapter 7	Process Integration and Electrification for Greenfield Design	171
7.1	Introduction.....	171
7.2	Methods	172
7.2.1	Base-Case definition, Electrification Target and Digital Twin Development .	174
7.2.2	Design Digital Twin and evaluation	174
7.2.3	Targeting (Pinch Analysis)	174
7.2.4	Targeting (Exergy Pinch Analysis).....	175
7.2.5	Net shaft-work calculation	181
7.2.6	Operating set-point optimisation.....	183
7.2.7	Process electrification technology selection and integration	184
7.2.8	Conventional and “Cross-Pocket” Heat Pump integration principles	185

7.2.9	Heat Exchanger Network design.....	186
7.2.10	Iterative design approach	187
7.3	Milk Evaporation Case Study Refresh	187
7.3.1	Solution constraints and ideals.....	189
7.4	Results.....	189
7.4.1	Initial Pinch Analysis, simulation, and heat pump integration results.....	190
7.4.2	Analysis of pressure set points for flash recovery and effect 2	192
7.4.3	Exergy Analysis for Design Improvement.....	194
7.4.4	Additional heat pump selection and placement	196
7.4.5	Heat exchanger network optimisation.....	197
7.4.6	Evaluation of design iterations.....	198
7.4.7	Final process design	200
7.4.8	Comparison with conventional heat pump integration	201
7.4.9	Evaluation of additional operating parameters and sensitivity analysis	202
7.5	Discussion.....	203
7.6	Conclusion	204
Chapter 8	Process Integration and Electrification for Retrofit Design.....	206
8.1	Introduction.....	206
8.2	Method.....	207
8.2.1	Preparation steps	207
8.2.2	Identification of potential retrofit opportunities.....	208
8.2.3	Expansion of the HSDT Method to include processes	210
8.2.4	Identify waste heat to “unlock” higher temperatures.....	211
8.2.5	Upgrading waste heat.....	212
8.2.6	Calculate levelised cost of heating.....	213
8.2.7	Iterative retrofit evaluation.....	214
8.3	Case studies	216

8.4	Results and discussion	218
8.4.1	Simulation and expansion of the HSDT method	218
8.4.2	Identify waste heat and application of conventional bridge analysis.....	219
8.4.3	Integration of heat pump technology	221
8.4.4	Estimate the impact and iterate	222
8.4.5	General retrofit solutions	223
8.4.6	Comparison with conventional retrofit solutions.....	226
8.4.7	Economic assessment.....	226
8.4.8	Other applications and limitations	228
8.5	Conclusions.....	229
Chapter 9	Conclusions.....	229
9.1	Progress to Implementation	231
9.2	Limitations and Future Work.....	232
Chapter 10	References.....	237
Chapter 11	Appendix.....	258
11.1	Appendix One: Code file snippets created for the Helmholtz formulation of Ammonia, n-butane and iso-butane.....	258
11.2	Appendix Two: Code snippets for surrogate property package modelling	285
11.3	Appendix Three: Refrigerant Efficiency Studies	288
11.3.1	The model assumptions.....	288
11.3.2	Case one: Efficiency analysis for varying sink outlet temperatures with a fixed sink inlet temperature of 15°C	289
11.3.3	Case Two: Efficiency analysis for the case of no subcooling for varying sink outlet temperatures with a fixed sink inlet temperature of 15°C	290
11.3.4	Case Three: Pressure considerations for the different refrigerants	291
11.3.5	Case Four: Efficiency analysis for varying inlet water temperatures with a fixed water outlet condition of 85°C	292
11.3.6	Case Five: Simple life cycle analysis (LCA) of direct and indirect emissions for R32.	294

11.3.7	Case Six: Summary of relevant safety standards	294
11.4	Appendix Four: All Heat Surplus-Deficit Tables for electrification retrofit.....	296
11.5	Appendix Five: Evidence of contributions to published articles.....	314

Table of Figures

Figure 1 The four electrification scenarios by Bühler et al. (2019).....	13
Figure 2 Decomposition of a process into hot and cold stream data (March, 1998).....	18
Figure 3 Composite (left) and Grand Composite Curve (right) demonstrating the targeting step (March, 1998).....	19
Figure 4 Generating a minimum energy requirement design following the Pinch rules (March, 1998).	20
Figure 5 Example of a heat loop and heat pathway (March, 1998).	20
Figure 6 Example of a heat path being utilised to reduce and eliminate utility heat exchange (March, 1998).....	21
Figure 7 The GCC being used to appropriately place a heat pump (Schlosser et al., 2021).....	22
Figure 8 Overview of Stampfli et al. (2019) hybrid heat pump integration method.	24
Figure 9 (a) Dimensioning of a heat pump using the GCC and COP curves. (b) the resulting heat pump specified (Schlosser et al., 2019).....	24
Figure 10 Advanced Composite Curve showing no cross pinch heat transfer (Nordman & Berntsson, 2009).	27
Figure 11 Plots showing the feasible regions for heat integration (Kamel et al., 2017).....	28
Figure 12 Original Energy Transfer Diagram identifying a heat bridge (Bonhivers et al., 2017).	29
Figure 13 Developmental stages for creating a Modified Energy Transfer Diagram (M.R.W. Walmsley et al., 2017).....	29
Figure 14 The Heat Surplus deficit table being used to identify good bridge pathways (Schlosser et al., 2021).....	30
Figure 15 Classifying EDT families using the proposed classification framework figure from (Yu et al., 2022).....	33
Figure 16 Microservice architecture deployed on Kubernetes sourced from (Burroughs et al., 2025)	36
Figure 17 Deviations in the representation of the sound speed data of Takagi (Takagi, 1996) for R-125 (figure from (Huber et al., 2022))	45
Figure 18 Process Integration and Electrification (PI&E) method.	61
Figure 19 Process Integration and Electrification retrofit design method through incremental changes.....	64

Figure 20 Base-case industrial milk evaporator integration illustrated using a process flow diagram.....	73
Figure 21 Comparison of Specific Heat Capacity (J/g·K) versus Temperature (°C) for whole milk (Composition of Fat equating to 4% by mass) against the model generated by Minim et al.(2002). Experimental data sourced from Lewis (1993).	90
Figure 22 Comparison of Specific Heat Capacity (J/g·K) versus Temperature (°C) for concentrated milk (Composition of Fat equating to 10% by mass). Experimental data sourced from Lewis (1993).	90
Figure 23 Comparison of Boiling Point Elevation (°C) versus Dry Mass (%) for a starting concentration of 90.7% water, 0.3% fat, 3.5% protein, 4.8% lactose, 0.7% minerals against the model generated by Winchester (2000) . Experimental data sourced from Winchester (2000).	91
Figure 24 Pressure-Enthalpy relationship data derived from for Ammonia Gao et al. (2023)	105
Figure 25 Pressure-Enthalpy relationship produced from generated package for Ammonia ..	105
Figure 26 Temperature-Entropy relationship data derived from Gao et al. (2023) for Ammonia	106
Figure 27 Temperature-Entropy relationship produced from the generated package for Ammonia	106
Figure 28 Error in Enthalpy Calculation vs Temperature and Pressure for Ammonia.....	107
Figure 29 Error in Entropy Calculation vs Temperature and Pressure for Ammonia	108
Figure 30 Error in Density Calculation vs Temperature and Pressure for Ammonia.....	108
Figure 31 Parity plots showing how a polynomial basis function fails to model the nonlinear system even with a significant number of training points.	118
Figure 32 Parity plots for RBF and Kriging surrogates for h_gas	119
Figure 33 RBF model for 100 data points of the humid air system.....	121
Figure 34 Kriging model for 100 data points of the humid air system.....	121
Figure 35 Parity plots to compare adaptive sampling for RBF and Kriging surrogates.....	125
Figure 36 Percentage difference in h_gas for the RBF surrogate model	126
Figure 37 Percentage difference in h_gas for the Kriging surrogate model	126
Figure 38 Code excerpt showing the injection of the surrogate and Helmholtz model	128
Figure 39 The general schematic of a TVR system sourced from Friso (2022).....	134

Figure 40 Experimental relationship between the diffuser efficiency and Mach number of the expanded motive gas sourced from Friso (2022).....	136
Figure 41 Enthalpy comparison between milk and water showing deviation	142
Figure 42 Code snippet showing the enthalpy offset term.....	142
Figure 43 Enthalpy comparison between milk and water showing deviation	143
Figure 44 DCH Model	144
Figure 45 Representation of the wetting nozzle air represented by large dark arrows with water being sprayed	145
Figure 46 Schematic of the inlets and outlets for a milk evaporator	146
Figure 47 Process flow diagram of an internal heat exchange with sub cooling cycle	148
Figure 48 A CO ₂ cycle exhibiting multiple pinch points throughout the length of the heat exchanger	150
Figure 49 Excerpt from the heat pump model showing the steep learning curve required for IDAES.....	157
Figure 50 Flowsheet of one of the cases to create an electric milk evaporator in Aspen Hysys	161
Figure 51 Flowsheet of one of the cases to create an electric milk evaporator in DWSim	162
Figure 52 Flowsheet of one of the cases to create an electric milk evaporator in the Ahuora Software Platform Powered by IDAES	163
Figure 53 Automatic Spreadsheet Data Collection (Hysys).....	165
Figure 54 Automatic Spreadsheet Data Collection (Python IDAES).....	166
Figure 55 Prototype for Automatic Data Extraction combined with Targeting	167
Figure 56 Process Integration and Electrification (PI&E) method.	173
Figure 57 Representation of Shifting Stream Temperatures.....	177
Figure 58 Representation of Pocket Cutting Strategies	179
Figure 59 Representative Net Heat Load Curve	179
Figure 60 Representative Net Exergy Load Curve	181
Figure 61 (a–c) Traditional air source heat pump with low COP overall, (d–f) Utilisation of a heat pocket to produce a higher COP.....	186
Figure 62 The base case design of a two-effect evaporator with MVR and TVR on each effect.	188

Figure 63 GCC of evaporation system illustrating a design with no vapour recompression (a), and with MVR integrated on Effect 1(b).	190
Figure 64 Option 2 a Double MVR System (No Effect 2 Vapour Recycling)	191
Figure 65 Development of an electric evaporation system (with integrated MVRs on Effect 1 and 2) using the GCC. (a) Non-optimised system design with single DCH. (b) Optimised design with single DCH.	191
Figure 66 Design with the same tube-side pressure on Effects 1 and 2, and optimised double DCH recovery.	194
Figure 67 GCC of MVR/MVR base case design with illustration of potential heat pump placement.	197
Figure 68 Heat Exchanger Network diagram of the electric milk evaporator.	198
Figure 69 Final process flowsheet design of the fully electric milk evaporation system.	200
Figure 70 Effect of changing the initial concentration, because of a preconcentration reverse osmosis operation, on total power consumption.	203
Figure 71 The relationship between the Exchanger Shifted (a) Composite Curve (ESCC), (b) EGCC, and (c) E-PTA for recovery exchanger as shown in Walmsley et al (2017).....	209
Figure 72 Process Integration and Electrification retrofit design method through iterative and incremental changes.	215
Figure 73 Case 1 – two-effect MVR-TVR effect evaporation system.	216
Figure 74 Case 2 – three effect MVR-TVR evaporation system.	217
Figure 75 Case 3 – four effect TVR evaporation system.	217
Figure 76 System boundary analysis of the evaporator (Case One).....	219
Figure 77 Example Heat Pump Bridge Analysis for the two effect MVR-TVR Case.	220
Figure 78 Example of DSI replacement for the two effect MVR-TVR case.....	224
Figure 79 Example of TVR replacement solution for the two effect MVR-TVR case.	225
Figure 80 LCOH of the three case studies for the TVR and DSI replacement design considering capital cost factors for planning, installation and infrastructure between $f_{C,Inst} = 2$ (solid lines) and $f_{C,Inst} = 4$ (dashed lines).....	228
Figure 81 COP at varying outlet sink temperatures for a 15°C water sink inlet.....	289
Figure 82 T-H Curve for R32 (left) and ammonia (right) showing the limited subcooling available for ammonia.....	290

Figure 83 COP at varying outlet sink temperatures for a 15°C water sink inlet without subcooling.	290
Figure 84 Maximum Cycle pressure	291
Figure 85 Pressure ratios for refrigerants at varying water outlet temperatures	292
Figure 86 Efficiency for different refrigerants with varying water inlet temperatures for a fixed water outlet temperature of 85°C	293
Figure 87 Phase envelopes for R32 at an inlet temperature of 55°C (left), and CO2 at a water inlet temperature of 45°C (right).....	293
Figure 88 DSI steam replacement retrofit for two-effect MVR-TVR design.....	296
Figure 89 DSI steam replacement retrofit for two-effect MVR-TVR design.....	297
Figure 90 DSI steam replacement retrofit for two-effect MVR-TVR design.....	298
Figure 91 TVR replacement retrofit for two-effect MVR-TVR design.....	299
Figure 92 TVR replacement retrofit for two-effect MVR-TVR design.....	300
Figure 93 TVR replacement retrofit for two-effect MVR-TVR design.....	301
Figure 94 DSI steam replacement retrofit for three-effect MVR-TVR design.....	302
Figure 95 DSI steam replacement retrofit for three-effect MVR-TVR design.....	303
Figure 96 DSI steam replacement retrofit for three-effect MVR-TVR design.....	304
Figure 97 TVR replacement retrofit for three-effect MVR-TVR design.....	305
Figure 98 TVR replacement retrofit for three-effect MVR-TVR design.....	306
Figure 99 TVR replacement retrofit for three-effect MVR-TVR design.....	307
Figure 100 DSI steam replacement retrofit for four-effect TVR design.....	308
Figure 101 DSI steam replacement retrofit for four-effect TVR design.....	309
Figure 102 DSI steam replacement retrofit for four-effect TVR design.....	310
Figure 103 TVR replacement retrofit for four-effect TVR design	311
Figure 104 TVR replacement retrofit for four-effect TVR design	312
Figure 105 TVR replacement retrofit for four-effect TVR design	313

List of Tables

Table 1 Fully commercialised mass manufacturable heat pumps that are designed to produce hot water (Poulsen & Zühlsdorf, 2022) (articles from Annex 58 describing the models are hyperlinked to the model name).....	17
Table 2 List of fluids translated into the equation-orientated Helmholtz formulations; uses and references	47
Table 3 Base case site stream data.	75
Table 4 Base case site utility stream data	76
Table 5 Sensitivity analysis of key design variables.....	76
Table 6 Final milk component properties used for process simulation (Zhang et al., 2014)....	84
Table 7 Specific heat coefficient values for fat, protein, lactose and salt used for simulation of raw milk in Aspen HYSYS (Munir et al., 2016).....	84
Table 8 List of Binary Interaction Coefficients for simulation of milk	85
Table 9 Symbol for Equation 5	94
Table 10 Physical constants and characteristic properties of ammonia Gao et al. (2023).....	101
Table 11 Coefficients of the ideal equation of state for ammonia Gao et al. (2023).....	102
Table 12 Coefficients of the residual equation of state for ammonia Gao et al. (2023).....	102
Table 13 Coefficients for saturated density equations (Gao et al., 2023)	103
Table 14 Surrogate model statistics across a range of models and data-points	116
Table 15 Adaptive sampling model statistics across a range of models and data-points for h _{gas}	123
Table 16 Variables to be changed by the user	151
Table 17 COP correlations for standard vapour compression, trans-critical CO ₂ and steam generation heat pumps developed by Schlosser et al (2020a).....	154
Table 18 Comparison of Values from different frameworks	164
Table 19 Representation of Gathering Stream Data	177
Table 20 Heat and exergy Pinch targeting for differing cases.	193
Table 21 Power requirements for the different design cases.	199
Table 22 Sensitivity analysis of key design variables on total power consumption.....	202

Table 23 Economic Parameters for LCOH calculation	214
Table 24 Stream data for the system boundary analysis of the evaporator (Case One)	219
Table 25 Steam savings and electricity cost across the three case studies.	224
Table 26 Economic parameters for the LCOH calculation.	227
Table 27 Results for simple LCA of R32	294

Nomenclature

Abbreviations

ACC: Advanced Composite Curve

ADT: Ahuora Digital Twin

BPE: Boiling Point Elevation

COP: Coefficient of Performance

COW: Condensate of Whey

CSD: Carnahan-Starling-DeSantis

DCH: Direct Contact Heat

DSI; Direct Steam Injector

DT: Digital Twin

ECS: Extended Corresponding States

EDT: Energy Digital Twin

EES: Engineering Equation Solver

EGCC: Exchanger Grand Composite Curve

EPTA: Exchanger Problem Table Algorithm

EPA: Exergy Pinch Analysis

EoS: Equation of State

ETD: Energy Transfer Diagram

FDM: Finite Difference Methods

GCC: Grand Composite Curve

GWP: Global Warming Potential

HEN: Heat Exchanger Network

HSDT: Heat Surplus Deficit Table

HVAC: Heating, Ventilation and Air Conditioning

HX: Heat Exchanger

IB: Insight Based

LCOH: Levelised Cost of Heating

MBIE: Ministry of Business and Innovation and Employment

MBWR: Modified Benedict-Webb-Rubin

METD: Modified Energy Transfer Diagram

MINLP: Mixed Integer Non-Linear Problem

MP: Mathematical Programming

MVR: Mechanical Vapour Recompression

NZEF: Net-Zero Energy Factory

OEM: Original Equipment Manufacturer

OPEX: Operating Expenditure

PA: Pinch Analysis

PE: Process Electrification

PI: Process Integration

PSE: Process Systems Engineering

PT: Pinch Temperature

PTA: Problem Table Algorithm

PV: Photovoltaic

PI&E: Process integration and Electrification

RBF: Radial Basis Function

STEP: Stream Temperature vs Enthalpy Plot

TDF: Temperature Driving Force

UHT: Ultra-High Temperature

VLE: Vapour-Liquid Equilibrium

Symbols

ΔT_{lift} Temperature lift (K)

ΔT_{min} Minimum approach temperature ($^{\circ}C$ or K)

x_{wsat} Maximum water vapour content

h_{gas} Molar enthalpy of the vapour phase

s_{gas} Molar entropy of the vapour phase

v_{gas} Molar specific volume of the vapour phase

G Mass flow for TVR

ω	Velocity
η	Isentropic efficiency or just efficiency
h	Specific enthalpy mass
\dot{m}	Mass flow
C_p	Heat capacity constant pressure
cp	Mass specific heat capacity constant pressure
U	Overall heat transfer coefficient
A	Effective heat transfer area
ΔT_{LM}	Log-mean temperature difference
ΔT_{cont}	Individual minimum temperature difference contributions

Subscripts

i	For TVR modelling induced side
m	For TVR modelling motive side
1	For TVR modelling the motive fluid as it exits the nozzle
2	For TVR modelling the state at the entry of the mixing chamber
3	For TVR modelling represents conditions at the exit of the mixing section
4	For TVR modelling denotes the state at the diffuser outlet
m	For DSI milk stream
s	For DSI steam stream
o	For DSI mixed outlet stream

Units

<i>K</i>	Kelvin
m^3	Cubic metres
<i>W</i>	Watt
<i>Wh</i>	Watt hour
$^{\circ}C$	Degrees Celsius
%	Percentage
Pa	Pascals
g	Gram
j	Joule

Prefixes

<i>G</i>	Giga ($\times 10^9$)
<i>M</i>	Mega ($\times 10^6$)
<i>k</i>	Kilo ($\times 10^3$)

Chapter 1 Introduction

1.1 Background

Climate change, and the associated drive to decarbonise energy systems, have placed increasing pressure on industrial sectors to transition toward low-emissions technologies. One of the most difficult challenges in this transition is the decarbonisation of industrial process heat, which accounts for approximately 25% of global final energy consumption and around 38% of energy-related CO₂ emissions (IEA, 2024). Despite growing momentum in the electrification of transport and power sectors, industrial heat remains overwhelmingly fossil-fuelled, particularly in high-demand sectors such as food processing, chemicals, and pulp and paper.

Most industrial process heat is supplied using steam or hot water that is generated by combustion of coal, natural gas, or biomass. Globally, over 50% of this heat demand falls into the low- to medium-temperature range (<400 °C), with a significant share of this heat below 200 °C (IEA, 2018). This temperature range is particularly well suited to electrification since commercially available technologies such as industrial heat pumps and mechanical vapour recompression (MVR) can achieve these temperatures with high coefficients of performance (COP). These technologies enable the substitution of fossil-derived thermal utilities with low-emissions electricity – offering a dual benefit of improving energy efficiency and reducing carbon intensity. However, widespread deployment of such technologies in both green (new) and brownfield (existing) industrial settings remains limited. A key barrier lies in the difficulty of identifying integration opportunities and quantifying the associated trade-offs. Many industrial facilities were never designed with electrification in mind, having had cheap fossil fuels available at the time, and retrofitting these systems involves both technical and economic challenges. These include constraints around process operability, equipment sizing, thermal integration, and capital expenditure.

Conventionally, Process Integration (PI) offered a structured thermodynamic framework to address some of these issues. Since its emergence in the late 1970s, PI has played a central role in optimising energy use in industrial plants, most notably through Pinch Analysis (PA). PA provides systematic methods for identifying energy targets and designing heat exchanger networks (HENs) to increase heat recovery within a process. Historically, these methods have been widely applied to improve the efficiency in which thermal energy is used in industrial plants. Efforts to adapt PI for electrification have made some progress. An example of this is through heat pump integration using PA. However, there are still significant limitations with these techniques. Traditional PI methods typically assume that the goal is to minimise thermal energy demand, an objective which is well-suited to systems

relying on fossil-fuelled utilities. However, this objective is poorly aligned with the characteristics of electrically driven technologies like heat pumps. Heat pumps are fundamentally constrained by their operational temperature lift – the temperature difference between the heat source and the heat sink. An increasing temperature lift decreases the COP, with a steep gradient at lower temperature lifts and eventually levelling off at lifts between 40 K and 120 K, according to Schlosser et al. (2020a). This means that even moderately high lifts require disproportionately more electrical energy input per unit of useful heat output.

Conventional Pinch-based targeting methods are not equipped to consider the temperature lift/COP trade-off. Heat has been treated as an interchangeable quantity with no differentiation between thermal energy sourced from fossil fuels and that supplied by work-driven systems. While thermodynamically feasible, conventional targets may be suboptimal economically and operationally when the heat is supplied using electricity. Additionally, there are a lack of heat pump integration design methodologies. As a result, integration opportunities are often misidentified. This necessitates a fundamental shift in focus, from demand reduction to electrical utility reduction. New Zealand provides a particularly relevant context for exploring these challenges. The national electricity grid is over 80% renewable, yet fossil fuels remain the dominant source of industrial heat. The dairy sector alone accounts for more than 60% of fossil fuel use in food processing, primarily for steam generation in evaporation systems. These systems operate in a temperature range well suited to heat pumps and MVR, but adoption has been slow due to uncertainty around integration and feasibility.

A key problem with conventional PA is the disconnect between the targets, based on idealised assumptions and simplified representations of the process, and the real-world interactions within the plant. Often, integration decisions are made based on steady state conditions, heuristic assumptions and basic property packages which are represented in a simplified flowsheet. This then leads to oversimplified Pinch targets and HEN designs that fail to consider the additional flow on effects that exist in large scale processes, before heat pump integration can even be considered. Digitalisation, and particularly the development of design digital twins, offers a promising pathway to overcome many of the limitations associated with conventional PI methods. A digital twin is a data-driven replica of a physical process or system that enables rapid simulation, optimisation, and scenario analysis. Unlike traditional PA that rely on static, idealised models, digital twins can capture the complex, non-linear interactions between heat and work, and capture the temperature-dependent performance factors such as heat pump COP. A digital twin allows engineers to test the impact of integration decisions under realistic conditions, quantify trade-offs, and identify robust solutions that

align with both thermodynamic and process constraints. As a result, digitalisation could be not just a complementary tool but a critical enabler for realising the full potential of PI and electrification in industrial decarbonisation.

Despite the promise of digital twins and advanced modelling environments, engineers face substantial challenges when attempting to digitalise industrial plants. Most commercial and open-source modelling environments are heavily geared toward oil, gas, and chemical applications, and their built-in unit operations, workflows, and thermodynamic property packages reflect this bias. For example, food, dairy, and bioprocess industries rely on complex, multi-component fluids such as milk, whey, or broth, which are not represented in standard component libraries. As a result, engineers are left with two difficult options: oversimplify the system and accept inaccuracies or attempt to build custom property packages or unit operations from scratch – an effort that is time-consuming, error-prone, and often beyond the practical capability of industrial teams. This lack of accessible, domain-relevant property models remains one of the most critical bottlenecks in scaling up digital twins and unlocking its promised benefits for industrial electrification.

Creating accurate digital twins requires a detailed understanding of process physics, including phase equilibria, heat and mass transfer, and fluid properties, knowledge that is often not widespread outside specialised process engineering teams. At the same time, implementing these models in platforms such as IDAES, Aspen, or Modelica demands proficiency in either, or a combination of, programming, mathematical optimisation, and model structuring, skills that are uncommon among plant engineers and rarely part of traditional process engineering education. These challenges are further compounded in industries outside the petrochemical sector, where software support is often weak or entirely lacking. Therefore, a generalisable method that combines Process Integration and Electrification (PI&E), supported by advanced digital twins, is urgently needed to address the decarbonisation challenges faced by modern industrial sites. To be effective, such a method must be underpinned by open-source tools that simplify both the creation of digital twins and the practical application of PI&E. This would help eliminate the current barriers of expertise and cost that prevent many sites (particularly in non-traditional sectors) from adopting low-emissions technologies and achieving meaningful decarbonisation outcomes.

1.2 Thesis Aim

The aim of this thesis is to synthesise a comprehensive Process Integration and Electrification (PI&E) methodology that underpins a core design module within a process digital twin platform context. A

particular focus is placed on the integration of industrial heat pumps, recognising their potential to economically decarbonise low to medium temperature process heat when effectively integrated into existing thermal systems.

The methodology extends conventional PI techniques by incorporating exergy-based targeting, retrofit constraints, and techno-economic evaluation, enabled and accelerated by advanced open-sourced digital tools. Therefore, the development and application of the proposed PI&E methodology requires a structured digital environment capable of representing complex process systems, integrating thermodynamic models, and supporting iterative design and evaluation workflows.

In this context, one of the main digital twin platforms used as the basis for this thesis is the Ahuora Digital Twin Platform (ADT). This platform is an in-house software, built by Ahuora – Centre for Smart Energy Systems with Stephen Burroughs (software engineer) and Benjamin Lincoln (chemical and process engineer, i.e. the present PhD candidate) who have been student co-leaders since its first lines of code were written in July 2022 (Burroughs et al., 2025). Although the implementation scope has focused on the ADT, the novel PI&E methodology could equally apply to other process digital twins and simulators (including Aspen Plus, Aspen Hysys, DWSIM, Unisim, Petro-sim, gProms, etc.). Indeed, some simulations have also been completed in Aspen Hysys and DWSIM. The ADT has the advantage that it is an in-house development and built on open-source Python libraries – IDAES (Lee et al., 2021) and Pyomo (Bynum et al., 2021), and state of the art non-linear solver – IPOPT (Wächter & Biegler, 2006).

Although the digital twin developed in this thesis is used primarily to enable the proposed PI&E methodology, it is not limited to this application. The digital twin also functions as a stand-alone process model that can support simulation, scenario analysis, operational understanding, data extraction, and evaluation of alternative process configurations independently of PI&E.

To achieve this aim, the research pursues the following steps:

- **Digitalise** the process problem by establishing a structured representation of industrial heat systems, enabling integration with a design-oriented digital twin environment.
- **Formulate** accurate property packages that can address modelling of complex fluids such as milk, humid air and unique fluids within the processing industries.

- **Target electrification opportunities** by applying an exergy-based process integration framework to quantify heat recovery limits, identify temperature mismatches, and determine where electrification technologies (e.g., heat pumps) can be thermodynamically and operationally integrated within industrial heat systems.
- Design and evaluate electrification solutions through an iterative synthesis framework that integrates selected technologies into the process flowsheet, assesses system performance, retrofit feasibility, and techno-economic viability, and validates the proposed PI&E methodology.

1.3 Structure of the Thesis

This thesis is structured to develop and apply a comprehensive methodology for PI&E, supported by advanced digital twins, and demonstrated through case studies in the dairy sector. The chapters progress from contextual background and literature review through to methodology development, detailed digital twin construction, and application to both greenfield and brownfield (retrofit) designs.

Chapter 1 establishes the background and motivation for this research, highlighting the challenges of decarbonising low- to medium-temperature industrial heat and the specific relevance of the dairy sector in New Zealand. It defines the overall research aim and objectives and outlines the scope of the work.

Chapter 2 critically reviews existing approaches to process integration, exergy analysis, and electrification technologies. It examines limitations of legacy simulation tools and identifies research gaps, particularly in applying electrification to food and dairy industries. The chapter highlights the need for a methodology that combines digitalisation with PI&E to overcome these limitations.

Chapter 3 presents the overarching PI&E methodology that forms the core contribution of the thesis. The methodology integrates three elements: digitalisation, process integration, and electrification design and evaluation into a unified framework supported by digital twins. The chapter is structured in two parts: Part A describes the creation of a digital twin, covering preparation, thermophysical property model development, and system modelling; Part B outlines how the digital twin is applied to process integration and electrification, including operational optimisation, energy and exergy targeting, and the systematic integration of electrification technologies. Both greenfield and retrofit applications are considered, with particular attention to the incremental and constrained nature of

retrofit pathways. This chapter establishes the methodological foundation for the remainder of the thesis.

Chapter 4 introduces the case study and focuses on preparing to build a digital twin of the dairy evaporation process. It defines design intent, system boundaries, operational requirements, and data collection strategies, and discusses the selection of appropriate simulation platforms. The purpose of the chapter is providing relevant information as to the scope and specific requirements for the later thesis chapters.

Chapter 5 develops custom thermophysical property models required for accurate digital twins in the food and dairy context. It covers activity coefficient models for milk, Helmholtz-based formulations for refrigerants, and machine-learned surrogate models for complex mixtures using humid air as the case study.

Building on Chapters 4 and 5, Chapter 6 describes the construction of the complete digital twin. It details the modelling of unit operations including evaporators, vapour recompression systems, and heat pumps, and demonstrates how simulation data can be extracted to support PI&E analysis.

In Chapter 7, the PI&E methodology is applied to greenfield design of a milk evaporation system. Conventional and exergy-based pinch analysis, operating set-point optimisation, and electrification technology integration are used to develop efficient process configurations. The results demonstrate how the methodology supports systematic electrification in new plant designs.

The methodology is adapted in Chapter 8 to account for retrofit constraints such as existing heat exchanger networks, waste heat recovery opportunities, and economic viability. Multiple retrofit case studies are presented to illustrate the approach.

The final chapter, Chapter 9, amalgamates the findings of the thesis, reflecting on the methodological contributions, the insights gained from greenfield and retrofit studies, and the broader implications for industrial decarbonisation. Limitations and recommendations for future research and development are also discussed.

Note: AI tools were used to edit and proof this thesis.

Chapter 2 Literature Review

2.1 Introduction

The decarbonisation of industrial process heat is widely recognised as one of the most pressing challenges in achieving global emissions reduction targets. Industrial heating processes account for a substantial share of both final energy demand and CO₂ emissions, yet the electrification of this sector has lagged behind the advances in the power and transport sectors. This slow progress reflects a combination of technical, economic, and operational barriers, particularly in low- to medium-temperature industries such as food, dairy, and pulp and paper. Within this context, process integration and electrification (PI&E) has emerged as a promising strategy to reduce reliance on fossil fuels while improving thermal efficiency and enabling the substitution of renewable electricity.

The purpose of this literature review is threefold. First, it establishes the current state of knowledge regarding PI methods, electrification technologies, and digitalisation tools relevant to industrial decarbonisation. Second, it critically evaluates the limitations of existing approaches, particularly the reliance on legacy simulation tools and conventional Pinch-based energy targeting methods, which were developed for petrochemical industries and are poorly suited to the requirements of food and dairy systems. Third, it identifies the methodological and technological gaps that motivate the research undertaken in this thesis, including the need for platform-agnostic methods, accurate thermophysical property models for complex fluids, and practical integration strategies.

The review begins by situating industrial process heat within the wider energy and emissions landscape, with emphasis on its significance globally and in the New Zealand dairy industry (the key case study for the thesis). It then surveys the range of available and emerging electrification technologies highlighting both their technical potential and their integration challenges. Subsequent sections examine the development and application of PI techniques, from the origins of PA to recent hybrid methods that combine exergy analysis, simulation, and optimisation. The literature on retrofit methods is also reviewed, recognising that most decarbonisation efforts will be applied to existing industrial plants rather than greenfield facilities. Finally, the review considers the transformative role of digitalisation and digital twins in enabling more accurate modelling, integration, and decision-making for PI&E, while contrasting these capabilities with the limitations of legacy commercial simulators.

Through this analysis, the review identifies a persistent research gap: while the electrification potential of low- to medium-temperature processes is well established, current methods and tools lack the flexibility, accuracy, and accessibility required to translate this potential into widespread industrial practice.

2.2 Process Heat Electrification and Industrial Decarbonisation

Industry accounts for ~37% of total energy-related CO₂ emissions globally (IEA, 2022), reflecting the sector's heavy reliance on fossil fuels to generate process heat. This fact has been relatively consistent, with emissions from industrial processes only declining by 2.3% to the year 2025 (IEA, 2025). Much of this heat is delivered to the process as steam and hot water. Decarbonising this heat therefore represents one of the most urgent avenues for near-term emissions reduction. In New Zealand, the challenge is particularly pronounced. Industrial process heat contributes a large share of national emissions, and the dairy sector is especially carbon intensive: ~92% of dairy process heat is supplied by coal and natural gas, with coal alone providing ~54% (MBIE, 2019). Despite clear policy signals to reduce fossil fuel use, these realities have maintained a strong dependency on combustion-based utilities.

Successive New Zealand governments have recognised that process heat presents one of the largest cost-effective opportunities to reduce national energy consumption while transitioning away from fossil fuels (MBIE, 2016). Government advice has consistently emphasised that industries should first minimise process heat requirements, as doing so lowers the capital and operating costs associated with replacing fossil-fuelled systems. Once demand is reduced, a range of low-emission alternatives become viable. Options most frequently recommended include biomass-fired boilers and heaters, as well as electrification pathways such as industrial heat pumps, electrode boilers, direct electric heaters, and more specialised technologies such as pulsed electric field processing (EECA, 2020b).

Despite sustained recommendations from government agencies, the uptake of electric process heating technologies in New Zealand has been markedly slower than the adoption of biomass. A prominent example is Fonterra's \$12 million conversion of a 43 MW coal boiler at its Te Awamutu site in 2020, a project that proceeded with only limited external support, securing \$200,000 from the EECA technology demonstration fund (EECA, 2020a). Since then, Fonterra and other major processors have continued to prioritise biomass conversions, with no indication of slowing investment in this pathway (RNZ, 2021). In contrast, the adoption of electrification technologies has slowed, with most projects requiring significant proportions of external financial assistance, such as the ANZCO hot water heat

pump initiative, to become commercially viable (EECA, 2025). Companies frequently cite high equipment, electricity, distribution and network upgrade expenses, in addition to limited in-house technical expertise, as the primary reasons for favouring biomass solutions over electrification. The situation described is not unique to New Zealand. Similar barriers to electrification have been identified internationally. A comprehensive survey by the United States Department of Energy, which examined a range of industrial case studies, highlighted the key factors hindering the uptake of electric technologies (Deason et al., 2018). These included:

Technical challenges:

- Limited commercial availability of electrification technologies for certain applications.
- Lack of engineering expertise to design or appropriately integrate electrified process systems.
- Heterogeneity across industrial sectors, which complicates the development of standardised solutions.

Economic challenges:

- Higher fuel and operating costs, with electricity generally more expensive than fossil fuels, making the economics of electrification less favourable.
- Elevated capital costs associated with fuel switching, including the expense of new equipment, ancillary services such as line upgrades, and the integration of electric technologies into existing plants.
- High costs or strict constraints on electric infrastructure capacity.
- Low profit margins in many high-temperature industries, which constrain the ability to absorb higher technology costs.
- While renewable electricity may be low-cost, the additional expense of firming supply or enabling operational flexibility often increases production costs or disrupts conventional plant operation and design.

Policy / Social Challenges:

- Regulatory and policy environments that either inadvertently discourage electrification or incentivise alternative fuels.
- Industrial risk aversion towards adopting new or unproven technologies.

Despite these challenges, international research and development into industrial electrification is growing. Additionally, increasing government investment in renewable electricity generation and distribution is providing both the incentive and the infrastructure for advancing electrification technologies (Atuonwu & Tassou, 2021).

2.2.1 Process electrification case studies

The growing interest in industrial process electrification has led to a number of case studies exploring both direct and indirect integration pathways. Chen et al. (2019) investigated electrification strategies for a methanol production plant using conventional PA techniques. Two scenarios were evaluated. The first considered indirect electrification, involving significant process modification to lower the required operating temperature by producing hydrogen via electrolysis and employing an alternative reaction pathway. The second scenario explored direct electrification through the integration of heat pumps, with the Grand Composite Curve (GCC) applied to determine the potential for utility savings.

The study yielded two important findings. First, the integration of heat pumps was found to be constrained by the limited temperature lifts achievable between available heat sources and sinks. As a result, direct electrification via heat pumps was able to substitute only a small fraction of the fossil fuel utility compared to the process-modification pathway. This highlights that, while direct electrification can contribute to decarbonisation, its effectiveness is strongly dependent on temperature lift constraints and should therefore be considered alongside other options such as process redesign.

Second, the study underscored the economic implications of electrification. The indirect pathway increased electricity demand from 14 MW to 604 MW, requiring an electricity price of approximately €4/MWh (NZD \$6.61/MWh) for the solution to be competitive with the base case. This finding demonstrates that large-scale indirect electrification is only viable in regions with abundant and inexpensive renewable electricity. It also highlights the critical role of government policy and electricity market structures in supporting an industrial transition to electrification.

Son et al. (2022) further clarified the distinction between electrification pathways, referring to direct electrification as *Power-to-Chemical* (in the context of chemical industries) and indirect electrification as *Power-to-Heat*. Their work developed a methodology for the systematic integration of Power-to-Heat technologies, with an emphasis on appropriate technology selection. Using an oil refinery as a case study, they evaluated candidate technologies based on operating conditions, technological readiness levels, and energy efficiency. The analysis concluded that heat pumps represent the most effective and holistic option for low- to medium-temperature heat applications, provided that the temperature lift (ΔT_{lift}) between source and sink is not excessively large. For high-temperature processes, however, electric heaters were found to be more suitable replacements for conventional fired heaters. Importantly, the study highlighted the necessity of a strong carbon price or emissions tax to make electrification economically competitive, reinforcing the importance of supportive policy frameworks in driving adoption.

The concept of an “uneconomically large ΔT_{lift} ” is also highly context dependent. As noted by Schlosser et al. (2020a), whether a given ΔT_{lift} is viable depends not only on process integration opportunities but also on the local electricity-to-fuel price ratio. This means that electrification strategies that are cost-effective in one region may not be viable in another. For instance, New Zealand’s electricity-to-gas price ratio is currently around 3 (MBIE, 2022b), compared to approximately 3.5 in Europe. This more favourable ratio suggests that electrification technologies such as heat pumps may achieve competitiveness in New Zealand sooner than in regions with less advantageous electricity pricing structures.

Another important consideration in the design of process utility systems is the decision to adopt either a centralised or a decentralised configuration, as each approach carries distinct advantages and drawbacks. Centralised systems typically employ larger units with greater thermal mass, which can result in operational inflexibility and longer implementation lead times. However, they often achieve economies of scale, since fixed capital costs are incurred only once and are distributed across a larger capacity (Ahlqvist et al., 2022). In contrast, decentralised systems are inherently more flexible, as the separation of units provides additional degrees of operational control. This flexibility comes at the expense of higher capital costs, as multiple smaller units must be purchased and maintained, thereby reducing economies of scale (Ahlqvist et al., 2022). In retrofit contexts, centralisation does not necessarily lead to better economics, as upgrades to a centralised system may require extensive modifications across the entire plant, whereas decentralised solutions can often be targeted to localised process areas.

Walmsley et al. (2016) further emphasised the complexity of this trade-off in a Total Site Integration study aimed at developing a 100% renewable milk powder plant at a small scale. Their work examined scenarios involving centralised renewable systems such as geothermal and solar thermal supply, in combination with decentralised electrification technologies such as heat pumps. The results demonstrated that centralised and decentralised systems can be mutually beneficial when integrated together, with centralised units delivering efficiency gains through economies of scale and decentralised units providing the flexibility required to balance variability and improve operability.

The number of full electrification studies in the dairy industry is extremely limited, with most research focusing on incremental technological developments that can be retrofitted into existing systems. Examples include the application of reverse osmosis or microwave heating technologies, as reviewed by Atuonwu & Tassou (2021), and the replacement of fossil-fuelled steam generation with alternative sources such as electrode boilers, biomass boilers, or large-scale steam-producing heat pumps (Bellemo & Bergamini, 2022). In contrast, studies addressing complete electrification of milk powder production systems are rare. Bühler et al. (2019) is one of the few examples, presenting four integrated electrification strategies: (i) a centralised electric boiler and cooling unit; (ii) the same configuration combined with process optimisations (reduced evaporator effects, replacement of thermal vapour recompression with mechanical vapour recompression, and air dryer heat recovery); (iii) a centralised integrated heat pump system with process optimisations; and (iv) a decentralised integrated heat pump system.

The heat pump-focused strategies are particularly noteworthy. In the centralised configuration, six heat pumps supplied heat to the evaporator (up to 90 °C), while a seventh high-temperature unit provided hot utility at 210 °C. The decentralised scheme employed the same seven heat pumps alongside two additional electric heaters. Bühler et al. (2019) found that the centralised heat pump system delivered the highest primary energy savings and lowest exergy losses but also suffered from significantly reduced operability due to the strong interdependence of process steps on the shared utility network. This issue was identified as a key barrier to industrial uptake, particularly in retrofit scenarios, where such interdependencies increase installation complexity and capital costs. Conversely, while challenging from a design and operability perspective, the study also highlighted a major positive finding: heat pump technology offered the lowest operating costs among the electrification pathways considered.

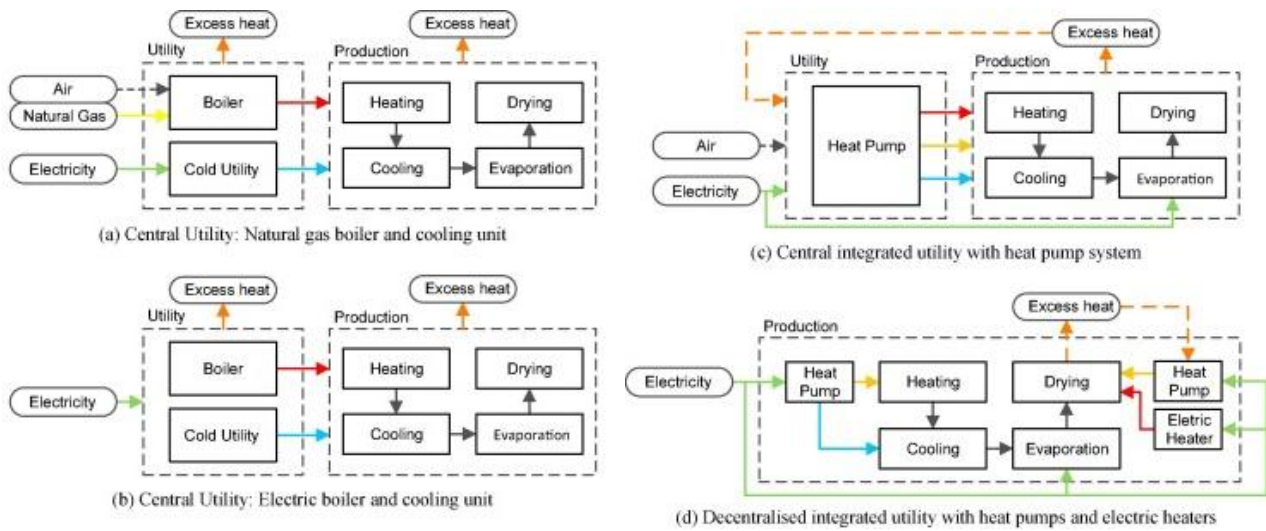


Figure 1 The four electrification scenarios by Bühler et al. (2019).

Wasser et al. (2023) present a detailed technical and economic analysis of a German cheese dairy redesigned as a Net-Zero Energy Factory (NZE) to assess both decarbonisation and grid flexibility potential. Using a simulation model in AnyLogic, the study integrates heat pumps, chillers, photovoltaic (PV) generation, and thermal/electrical storage to replace conventional gas boilers and refrigeration units. The electrification of thermal processes (particularly through a 324 kW heat pump system) was identified as the most impactful measure, reducing primary energy demand while providing substantial operational flexibility. The system demonstrated the ability to shift loads by rescheduling process chains (milk, yogurt, cheese), achieving reductions in power demand of up to 51.7% in some periods, while also accommodating load increases of 86.8% at other times. Economic analysis indicated that PV and heat pumps were cost-effective compared to fossil alternatives, though battery storage remained uneconomical without external remuneration.

Extrapolation of the results to the national scale suggests that converting all German dairies into NZEFs could yield daily flexibility potentials of -2.38 GWh to +2.5 GWh, with maximum power contributions of -540 MW to +930 MW. This would allow dairies not only to decarbonise their operations but also to act as active players in stabilising electricity grids with high shares of intermittent renewables. It concludes that NZEF transformation of dairy plants offers both environmental gains and an additional grid-balancing instrument, positioning the sector as a valuable contributor to Germany's wider energy transition.

While Bühler et al. (2019) and Wasser et al. (2023) provide two of the most comprehensive studies on electrification of dairy systems, both exhibit important limitations that restrict their applicability to large-scale milk powder production and the retrofit-heavy contexts that dominate industry practice.

Bühler et al. (2019) presented four electrification strategies for milk powder plants, including centralised and decentralised heat pump systems. While the study demonstrated the technical feasibility and relative efficiency of these approaches, its methodology was largely confined to conceptual energy targeting using simplified heat balances. The process models did not incorporate rigorous thermophysical property methods, nor did they capture the complex fluid behaviours of milk concentrates. This introduces uncertainty into both energy balances and technology integration outcomes. Furthermore, the centralised heat pump strategy, while efficient in principle, did not align with findings from other literature which emphasise that temperature lift should be minimised wherever possible to improve heat pump performance and economic viability. By relying on large, centralised high temperature, the study introduced significant work penalties that could have been mitigated by a more distributed or stepwise integration strategy.

In contrast, Wasser et al. (2023) adopted a NZEF framework, modelling a German cheese dairy with heat pumps, PV, and storage to provide grid flexibility. While this work advances the discussion of electrification beyond plant boundaries to include energy system integration, its scope is restricted to cheese production, omitting evaporation and spray drying (which is the most energy-intensive processes in milk powder manufacturing). As such, its findings underestimate the sector's decarbonisation challenge. In addition, the modelling approach relied on high-level simulation with simplified load-shifting heuristics, rather than rigorous process simulation capable of testing operability under varying product conditions.

Together, these studies illustrate the growing momentum towards full electrification but leave several gaps. Specifically, they lack: (i) in depth modelling of evaporation and drying systems, (ii) rigorous property modelling for complex fluids, where thermodynamic inaccuracy undermines process integration and technology evaluation; (iii) retrofit-focused integration strategies that balance work efficiency with operability; and (iv) demonstration of digitalisation methods such as digital twins or exergy-based targeting to validate electrification pathways in realistic, plant-specific contexts.

2.2.2 Case study in context: milk powder

Milk powder production is a critical component of the global dairy industry, offering a shelf-stable and transportable form of milk that supports food security and international trade. The process involves removing water from raw milk through thermal operations such as evaporation and spray drying, which are both energy intensive. These thermal processes typically rely on fossil-fuel-derived steam, with evaporation and drying together accounting for over 90% of the energy used in milk powder production (MBIE, 2022a). Globally, the milk powder market continues to grow, with an estimated value of \$59.39 billion NZD in 2024 and projected growth to \$ 89.95 billion NZD in 2025 (Future Market Insights, 2024). Major producers include the European Union, the United States, and New Zealand, where large-scale production facilities operate for significant portions of the year to convert fresh milk into powdered form. Given the substantial energy use and carbon footprint associated with these facilities, the milk powder sector has become a focal point for industrial decarbonisation and electrification efforts.

New Zealand plays a leading role in the global milk powder trade, consistently ranking among the top three exporters. In 2023, dairy exports reached NZD 18.2 billion, with milk powder representing a significant portion of this value (Dairy News Today, 2024). The industry processes more than 20 billion litres of raw milk annually, and large plants routinely operate at scales exceeding one million litres per day during peak season. As a result, milk powder production accounts for a major share of the country's industrial energy use. In 2016, the dairy processing sector in New Zealand consumed 28.4 petajoules (PJ) of fossil fuel energy to generate process heat, resulting in approximately 2.1 million tonnes of CO₂-equivalent emissions (MBIE, 2016). Coal and natural gas made up the majority of this energy input, with coal alone accounting for 54% of fuel use. This energy demand is equivalent to the annual consumption of around 800,000 New Zealand households, highlighting the scale of opportunity for electrification (MBIE, 2016).

Given the high thermal energy demand, heavy reliance on fossil fuels, and New Zealand's access to a largely renewable electricity grid, the milk powder application presents a compelling case for implementing PI&E strategies. The presented case study focuses specifically on the milk evaporation system, which is a thermally intensive sub-process within milk powder manufacturing. The selected system serves as a representative model to demonstrate the application of the methodology developed in this research.

2.2.3 Standardised commercially available heat pump review

Industrial large-scale heat pumps are a well-researched and established process heat technology, capable of providing low emission heat below 100°C (Adamson et al., 2022). The following review will focus on the already established and marketable heat pump technology that could supply medium temperature heat (60 – 120 °C). The section aims to determine what devices are available in New Zealand and internationally, whilst also determining common design philosophies that make the trade-off between capital costs and cycle performance more favourable. Medium temperature heat pumps currently available on the market are usually custom designed units that are designed for the specific applications (EECA, 2019). The equipment is generally sourced from large Original Equipment Manufacturer (OEM) refrigeration suppliers. This is a stark contrast to low temperature heat pump units' consumers commonly purchase, which are typically off the shelf units. Medium temperature heat pumps follow the same core principles all heat pumps utilise to operate. Absorbing waste or low-grade heat into a purpose selected refrigerant (the source), upgrading its temperature through a compression step, then distributing the upgraded heat into another medium (the sink), then finally expanding the refrigerant back to the initial conditions to repeat the cycle.

Commonly commercially available cycles often use ammonia (R717) for large scale industrial sites due to its low global warming potential (GWP) and high efficiency in the 70°C to 85°C temperature range (EECA, 2019). The downside of this refrigerant is its hazardous nature – being highly flammable and toxic at low concentrations, therefore requiring expensive safe management and mitigation procedures. Other common refrigerants used are butane (R600) and iso-butane (R600a) commonly found in dual refrigeration – heat pump units. Butane operates at a much lower pressure than many of the common refrigerants, however, has had its application in higher temperature uses limited by its flammability and explosion risk. Carbon dioxide (R744) is one of the oldest refrigerants still used in cycles today. CO₂ is most commonly available in transcritical cycles, due to its low critical temperature of 31°C. Instead of providing a constant temperature output like the condensing ammonia and butane cycles, the temperature will glide, which is similar to the characteristics of a heat exchanger using a super-heated or sub cooled fluid (Sarkar et al., 2005). This can provide efficiency benefits if the temperature profiles of the refrigerant and sink can be matched. CO₂ systems have limited applications due to the low evaporation temperatures and its extremely high operating pressures.

2.2.4 Commercially available mass manufactured heat pumps internationally

The European and Japanese heat pump markets are more advanced in maturity than in New Zealand (Poulsen & Zühlsdorf, 2022). Currently, these markets have produced a wider range of fully commercialised industrial heat pumps at with maximum temperatures and duties ranging from 120 °C - 212 °C and 0.03 MW – over 70 MW (high range values are stackable MVR units from Piller (2022)). These heat pumps have many different designs and configurations to achieve a wide coverage of process conditions and capacities. Table 1 details heat pumps at a TLR of 9 designed for medium temperature hot water production.

Table 1 Fully commercialised mass manufacturable heat pumps that are designed to produce hot water (Poulsen & Zühlsdorf, 2022) (articles from Annex 58 describing the models are hyperlinked to the model name)

Supplier	Compressor type	Working fluid	Capacity	Max supply temperature
Fuji Electric	Reciprocating	R-245fa	0.03 MW	120 °C
Mayekawa (EcoSirocco)	Reciprocating	R-744	0.1 MW	120 °C
Mayekawa (EcoCircuit)	Reciprocating	R-1234ze(E)	0.1 MW	120 °C
Kobelco Compressors Corp. (SGH120)	Twin-screw	R-245fa	0.4 MW	120 °C
Mitsubishi Heavy Industries	Two-stage centrifugal	R-134a	0.6 MW	130 °C
Combitherm	Semi-hermetic screw	R-1233zd(E)	0.3 – 3.3 MW	120 °C
Hybrid Energy	Piston, Screw	R-717, R-718	0.5 – 5.0 MW	120 °C
Heaten	Reciprocating, custom design	HFOs	1.0 – 6.0 MW	165 °C

Table 1 is not fully representative of all heat pumps available, as there are some promising designs progressing through the research and pilot stages. However, the aim of the review to use available technology to accelerate uptake of electric technologies and through that lens the heat pumps listed above are the heat pump designs that fulfil the listed criteria.

2.3 Process Integration for Energy Efficiency

Klemeš & Kravanja (2013) describe Process Integration (PI) as a “holistic approach to process design and operation that emphasises the utility of the process”. Their approach maximised efficiency, in traditionally fossil fuel-based systems, through increasing process heat recovery and minimising

utility demand, which subsequently leads to a reduction in CO₂ emissions (Linnhoff & Flower, 1978). Before 1970, PI was not a prevalent field due to low-cost energy sources, and a lack of understanding of the environmental impacts of using fossil fuels. However, during the 1970s oil crisis, the sustainable and economic management of energy resources became a greater focus (Klemeš & Kravanja, 2013). During this time, several methodologies began to emerge, with Process Systems Engineering (PSE) being the most prominent (Sargent, 1979). This method was more focused on flow sheeting and simulating process behaviour to generate energy efficient designs, but this method proved difficult for the user to determine if the generated process design was optimal. Another method produced was Pinch Analysis (PA) (Linnhoff & Flower, 1978), which has remained the core area of PI implementation and research for the past 40 plus years (Klemeš & Kravanja, 2013).

2.3.1 Pinch analysis

PA is comprised of four key concepts, decomposition of the process, targeting thermodynamic limits, design generation, and optimisation. PA decomposes the process into a series of hot and cold streams, as shown in Figure 2 (March, 1998). Hot streams are streams which require cooling and can “give heat” to other streams, whereas cold streams require heating and can to “receive heat” from other streams.

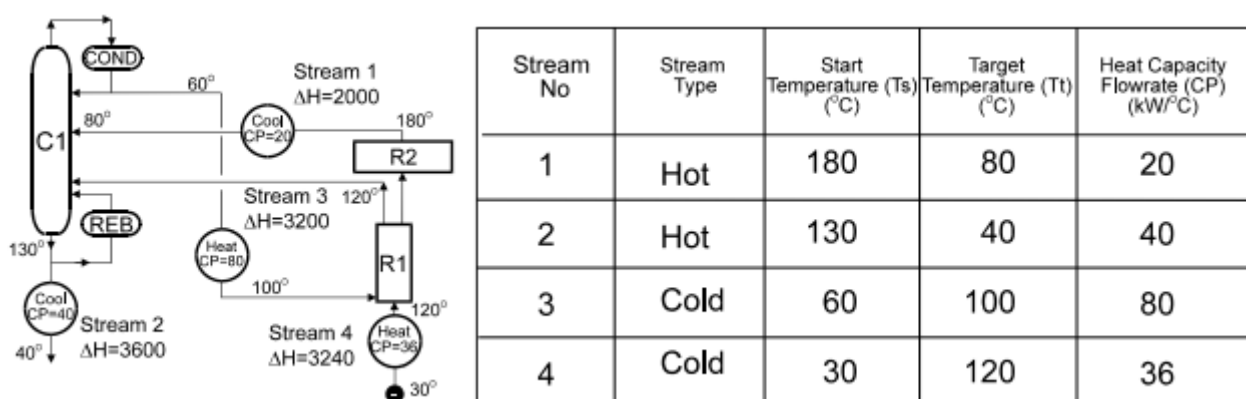


Figure 2 Decomposition of a process into hot and cold stream data (March, 1998).

A Problem Table Algorithm (PTA) is used to process the stream data and determine areas of heat surplus and deficit, which can be visualised as hot and cold composite curves (Figure 3). Composite curves represent how heat transfers or cascades through the system, subsequently indicating the minimum energy targets for hot and cold utility and maximum amount of heat recovery that are thermodynamically obtainable. Additionally, the PTA can also be used to produce the Grand

Composite Curve (GCC), also shown in Figure 3, which represents the net surpluses and deficits of heat at certain temperature levels. Understanding the GCC, and subsequently the process net heat cascade, allows for appropriate utility to process matching (March, 1998). Correct matching of the appropriate utilities improves the fidelity of the targeting step, elevating a simple MW target to multiple energy targets that considers utility quality and therefore a representative cost.




Figure 3 Composite (left) and Grand Composite Curve (right) demonstrating the targeting step (March, 1998).

After understanding the thermodynamic limits for PI, PA provides heuristics to assist solution generation. Two key features are the concepts of the Pinch Temperature (PT) and minimum approach temperature (ΔT_{\min}). The ΔT_{\min} is an economic constraint, chosen by the user, that limits how close in temperature a hot and cold stream can exchange to. A higher ΔT_{\min} will decrease the amount of heat recoverable, however it will also decrease the required heat exchange area, and vice versa. As the field of PA developed, typical ΔT_{\min} s for different applications became recognised. For example, 40 °C between flue gas and a process stream, and about 3°C for low temperature refrigeration (March, 1998). The PT is the point where the hot and cold surpluses converge to the ΔT_{\min} . The PT is significant, as it denotes where the process has a surplus of heat (above the PT) and where there is a deficit of heat (below the PT). This concept forms the three fundamental rules of HEN design: 1) do not supply hot utility below the PT (because there is already a surplus of low-grade heat), 2) do not supply cold utility above the PT (because there is already a deficit of heat), and 3) do not recover heat across the PT (because heat will be transfer from a portion of heat deficit to a portion of heat surplus). The three fundamental rules assist in designing an energy efficient HEN/process. The thermodynamic targets also provide a good measure to the effectiveness of the HEN design. An example of a HEN design by March (1998) is shown in Figure 4.

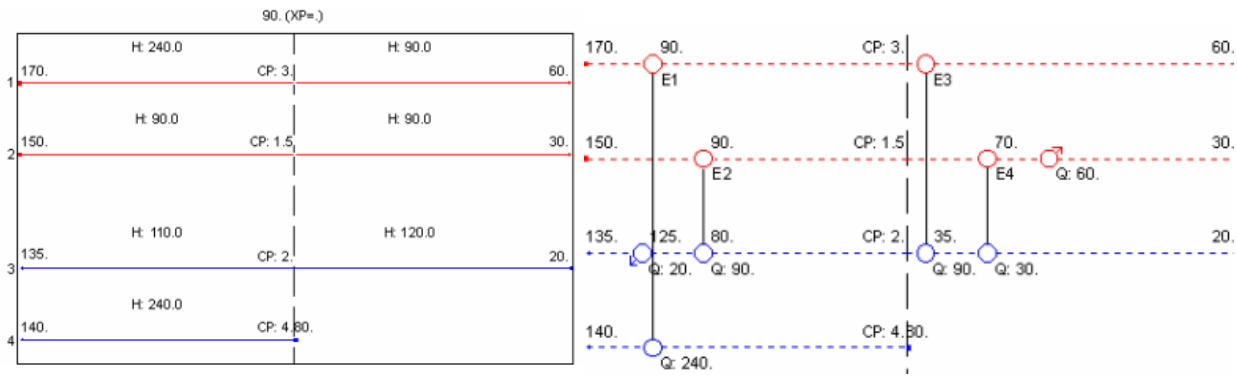


Figure 4 Generating a minimum energy requirement design following the Pinch rules (March, 1998).

The final step of conventional PA is the optimisation/evolution of the network to meet the user's criteria. A minimum energy design pathway can cause the resulting HEN to be overly complex (containing a large proportion of operations such as splits and small heat exchangers (HXs)). These designs can be capitially expensive and potentially difficult to control; therefore, the HENs design can be modified to reduce complexity without a significant loss of energy recovery. The main Pinch method for network evolution is utilisation of heat pathways and loops (Figure 5).




Figure 5 Example of a heat loop and heat pathway (March, 1998).

A heat loop is a circuit through multiple HXs, where the internal duty of the circuit can be shifted around the circuit without affecting the other HX duties external to the loop (March, 1998). This phenomenon can be leveraged to remove small HXs by shifting the small HXs duty onto the larger exchanger in the loop, hence removing complexity with minimal effect to the HENs performance.

A heat path is a continuous path connecting a heater to a cooler through the network. Because of this connection, if a charge to the heat load of a recovery exchange is made the heater and cooler load will change by the same amount (but in the opposite direction). As a result, a utility HX can be eliminated or reduced by increasing the recovery exchangers load in the heat path (Figure 6).

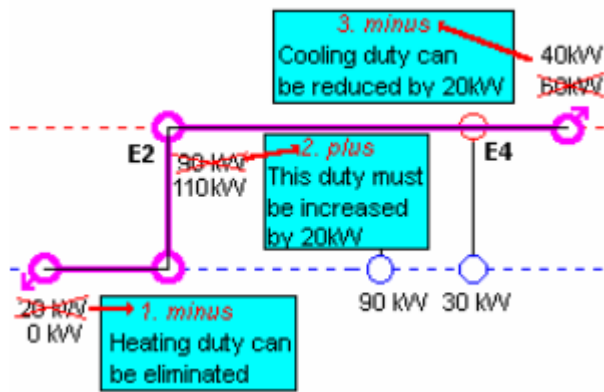


Figure 6 Example of a heat path being utilised to reduce and eliminate utility heat exchange (March, 1998).

2.3.2 Idealised Pinch Analysis heat pump integration methods

The first mention of heat pump integration in the literature was from Townsend & Linnhoff (1983) where they extended their original Pinch concepts to include power systems. The paper specially outlined the criteria for appropriately integrating a heat pump, a principle that has been observed to present day. Townsend & Linnhoff (1983) said that when attempting to integrate a heat pump the most effective place is across the PT, with the evaporator of the heat pump below the PT and the condenser above the PT. A heat pump placed here would take heat for areas of surplus, upgrade it, then deposit the upgraded heat in areas of heat deficit. If this principle was not followed, then the fundamental Pinch rules would be violated, and inefficiencies would occur. It then describes that the GCC is an effective tool to target the optimal placement of said heat pump (Figure 7)

Wallin et al. (1990) expanded on the work of Townsend & Linnhoff (1983) to allow a user to determine which heat pump type was optimal. This was achieved by matching the shape of the process composite curve against the composite curve of the specific heat pump option. The closer that the two curves resembled each other the better the specific type of heat pump would operate for that given process. The paper also initiated the research into optimizing a heat pump against annualised capital costs, via manipulation of heat source and sink temperatures and, subsequently, heat pump duty.




Figure 7 The GCC being used to appropriately place a heat pump (Schlosser et al., 2021).

In a generalised sense, integration methods are categorised into two themes, Insight Based (IB) and Mathematical Programming based methods (Sreepathi & Rangaiah, 2014). In more recent years hybrid methods that combine elements of IB and MP have also been popularised. This fact holds true for heat pump centred integration methods. The two papers from Townsend & Linnhoff (1983) and Wallin et al. (1990) founded both IB and MP heat pump integration fields with more recent methods building from these base founding principles.

2.3.3 Current heat pump integration methods

Schlosser et al. (2019) published a review of PA for heat pump integration in an industrial context. The review aimed to demonstrate the energy savings potential of PA for heat pump integration as well as updated the method to utilise modern tools, expressions and databases. The paper integrated the conventional PA method with the MP-based methods developed by Stampfli et al. (2019). Both papers used a hybrid approach for heat pump integration, with Stampfli et al. (2019) adding the additional complexity of non-continuous processes. Both papers selected a hybrid approach (of IB and MP) to cover the weakness of each individual method. The use of IB techniques helped to ensure the design solution generated was both practical and applicable, something which MP methods can struggle to do. Also, the inclusion of IB techniques allowed for the search space for MP to be dramatically reduced, allowing for less computational power being required. The MP portion of the

paper covered the weakness of IB techniques, including being slow, cumbersome and impossible to guarantee optimality in large complex integration problems (Stampfli et al., 2019).

Figure 8 gives an overview of the method used in Stampfli et al. (2019). It starts by extracting process data and identifying areas of heating and cooling requirements, as well as scheduling data for non-continuous processes. Next, the targeting stage uses PA and the corresponding PA tools (such as the GCC and PT) which are combined with a COP curve to optimally model the evaporation temperature and correlating evaporation duty (the COP curve is the dotted blue line in Figure 8) of the chosen heat pumps. Stampfli et al. (2019) defines the COP curve as the absorbed heat flow of the HP as a function of the temperature levels of condensation and evaporation. Becker (2012) suggested that, for current heat pump devices, an appropriate 2nd law efficiency of 0.55 would be representative. This value was verified by Arpagaus et al. (2018), who conducted experimental data collection on a wide range of industrial heat pumps, finding that a range of between 0.4 to 0.6 was typical.

The COP curve can be plotted on the GCC of the process (Schlosser et al., 2019), as shown in Figure 9. The condensation and evaporation temperatures are then the basis for the MP optimisation later in the method. The MP objective function is made up of annual investment costs, network operational costs and environmental footprint costs. The MP then optimises the initial network produced via the IB (PA) method to reduce cost and emissions into the optimal solution. Schlosser et al. (2019) found that this combined integration technique proved to be highly effective, with one case study reducing operating costs by 61% and the emissions by 52% when compared to the original system.

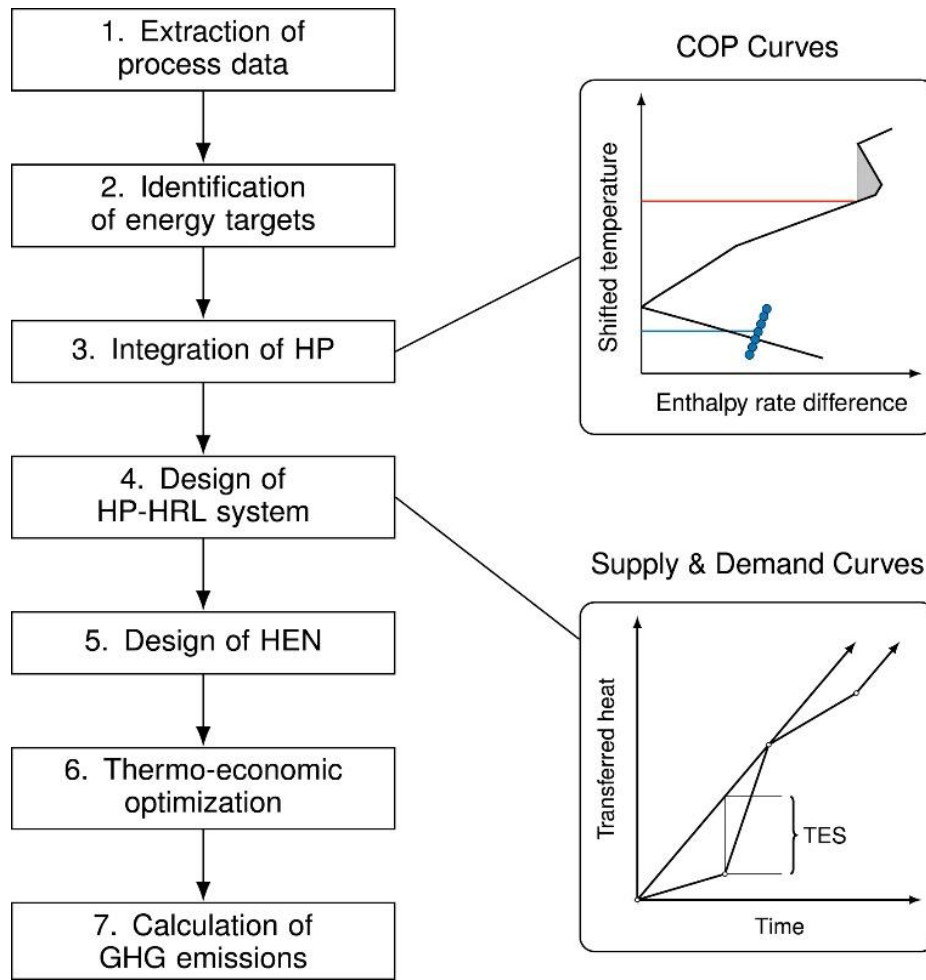


Figure 8 Overview of Stampfli et al. (2019) hybrid heat pump integration method.

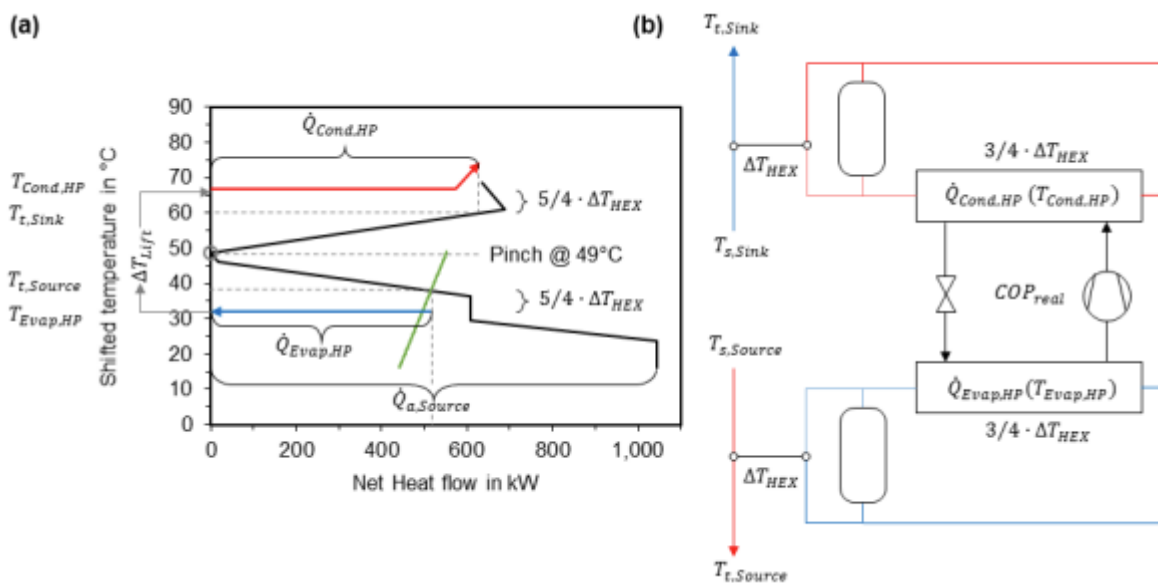


Figure 9 (a) Dimensioning of a heat pump using the GCC and COP curves. (b) the resulting heat pump specified (Schlosser et al., 2019).

2.3.4 The exergy per total annual cost method (EPC)

COP is a common selection metric for heat pump integration, however, Wang Mengying et al. (2018) proposes using exergy as an alternate evaluation and selection tool. They argue that different heat pump systems utilise multiple types of energy quality, for example, electricity and steams at differing pressures, therefore, a simple metric like COP does not fully capture and evaluate the effectiveness of a selected heat pump. In simple terms, exergy is a measure of energy that is available for use. For example, water that is at atmospheric conditions on the ground has zero exergy since no useful work can be generated from it (Wall, 1986). However, steam at 300°C (saturated) has an abundance of exergy. The other important concept is exergy losses. As a material transforms and there is a conversion of energy into different forms, some entropy will be generated which lowers the available exergy of the system as useful work is transformed into lower grade heat that becomes no longer useful. Wang Mengying et al. (2018) takes the exergy concept further and puts it in terms of exergy per total annual cost.

In the paper, the exergy metric is used in a similar way to most other PA based methods. The GCC is used to determine where a potential heat pump would be integrated, then different heat pump technologies are placed in the suggested locations and evaluated against the proposed criteria. After the appropriate heat pump is selected (the one that gives the highest exergy per annual cost), the method concludes by using the conventional PA method for HEN design to configure the rest of the system. The method presented in this paper highlights that it is difficult to target and determine whether your design is achieving minimal work or cost. However, the exergy per annual cost does not provide any additional insight to whether the design is approaching optimal conditions.

2.3.5 Simulation based methods

Yang et al. (2020) proposed a method for combining PA with rigorous process simulation. The work first starts with the conventional PA based method for integrating a heat pump but also utilises process modelling software: Aspen HYSYS and MATLAB. The real process design is converted into a digital model. The digital model is then fed into an MP based optimisation to minimise total energy consumption. Yang et al. (2020) then proposes a meta-heuristic genetic algorithm approach to solve the MP model since the problem lacks a smooth continuous gradient making it challenging for linear solvers. The method takes a variety of potential solutions, simulates them, determines the total energy usage, pairs the best solutions to mate, add random mutations and repeat till the most economic heat

pump placement is found. The method was evaluated using a case study and the design it produced achieved an energy saving of up to 50% when compared against the unintegrated base case showing a simulation-based approach to be favourable.

2.3.6 Integration for non-minimal Pinch target systems

A common assumption made by all approaches discussed thus far is that the HEN, in which the heat pump is integrated, will meet the minimum pinch targets. This assumption is made by using the GCC to appropriately place the chosen heat pump, as the GCCs shape is only achieved by creating a minimum pinch target solution. Schlosser et al. (2021) noted that real plants find it extremely challenging from a technical, spatial and cost perspective to achieve this assumption, therefore, it's most likely that some Pinch violations will occur in real systems. Designing for the idealised scenario could potentially miss additional demand requirements, in turn reducing the heat pump efficiency and sizing as the whole picture is not fully understood. Schlosser et al. (2021) states that this is currently a major problem in the space of heat pump integration as a whole (especially in retrofit) and that their paper is one of the only that address the issue directly.

2.4 Heat Pump Retrofit

Retrofit solutions are the key to accelerating the decarbonisation of the process heat sector since most decarbonisation will need to occur in existing plants. However, retrofit methods are lacking in the open literature, compared to heat pump integration methods for greenfield designs. In part, that is due to the ability to adapt, to some extent, a greenfield integration method for a retrofit context. However, since many IB methods are based on the minimisation assumptions within PA and there are many solutions that could be generated through MP methods, the additional complexity and variables involved in retrofit brings can cause non-specialised methods to find non-optimal solutions. This section will cover the methods that do exist for heat pump retrofit. To establish the foundations for these methods, brief mention will also be made for HEN retrofit methods.

2.4.1 Heat Exchanger Network (HEN) retrofit methods

Normann & Berntsson (2009) put forward the concept of the Advanced Composite Curve (ACC) to show where cross pinch violations occur (Figure 10), which was a significant challenge in previous HEN methods. This was found to work well for small problems, however as the number of streams

and HX units increased, the method became confusing and cumbersome, also proving challenging to automate.




Figure 10 Advanced Composite Curve showing no cross pinch heat transfer (Nordman & Berntsson, 2009).

Individual Stream Temperature vs Enthalpy Plot (STEP) also aimed to identify cross pinch violations (Y.Q. Lai et al., 2017). The major idea of STEP is to match the process streams directly to the heat recovery plot, therefore allowing the user to have the layout of the existing network on the diagram. Subsequently, this highlights cross pinch transfer and validates if new HXs maintain temperature and process constraints. However, the STEP plot also becomes confusing as the number of streams increase. Schlosser et al. (2021) also provided critique on STEP, stating that the method does not calculate any form of energy targeting, therefore making it difficult to prove the optimality of the retrofit solution.

Other methods have attempted to move away from Pinch composite curves. The Temperature Driving Force diagram (TDF) plots the temperature driving forces for all HXs within a process (Kamel et al., 2017). In a TDF, the HXs are represented as straight lines (Figure 11). The chart then plots the hot and cold pinch temperatures and then segregate the plot into five regions. Each region is classified as either a good area of heat recovery or a poor area, such as HXs crossing any Pinch lines.

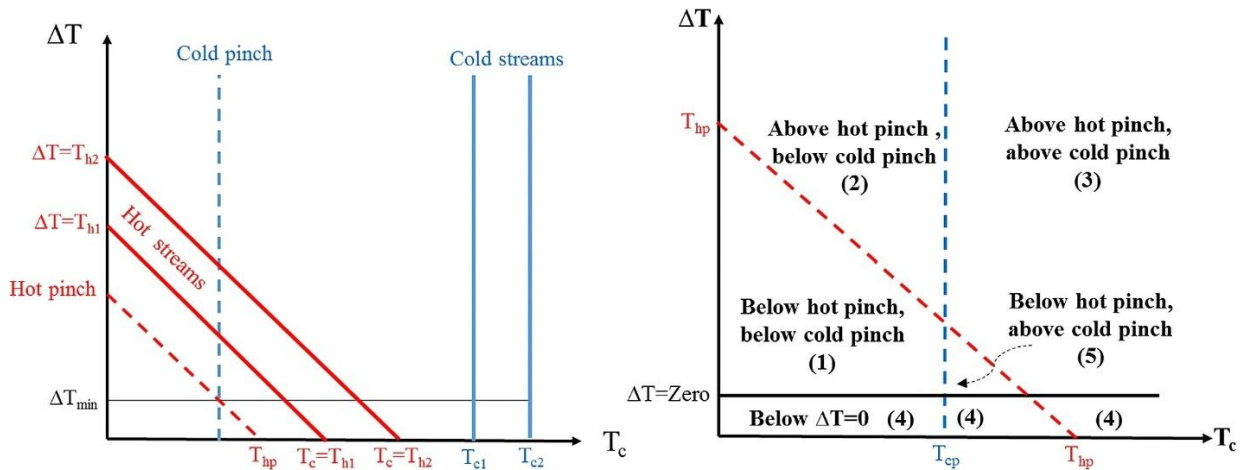


Figure 11 Plots showing the feasible regions for heat integration (Kamel et al., 2017).

2.4.2 Modified Heat Pump Bridge Analysis

The method of Bridge Analysis was first incepted by Bonhivers et al. (2017), formulating the idea of using the Energy Transfer Diagram (ETD) to find energy saving pathways (bridges) between heaters and coolers. These bridges could then be utilised to convey excess heat from a cooler to an area where there was not enough heat, i.e. where there is an existing heater (Figure 12). The ETD also helps to identify cross Pinch violations and areas of high potential temperature driving force (combining the ideas from other previously discussed methods).

Walmsley et al. (2017) noted that the overall concept of the ETD was useful, however, the format of the ETD was unintuitive. In this paper, the authors proposed a more recognisable format which was named the Modified Energy Transfer Diagram (METD). The METD represents a “shifted GCC”, where the conventional GCC was shifted to the right, matching the deviation for the thermodynamic minimum solution (the 1950 kW value shown in Figure 13d).

Lal et al. (2018) further developed the bridge method by proposing a more numerical tabular version called the Heat Surplus Deficit Table (HSDT). The table allowed for many more streams to be included as the graphical representation of the METD becomes too cluttered with a large number of streams (an issue faced by most IB methods). Additionally, a benefit was introduced that allowed the quantification of retrofit savings to be determined easily, as the numbers were already available. The method was also able to be automated (N. S. Lal, 2020).

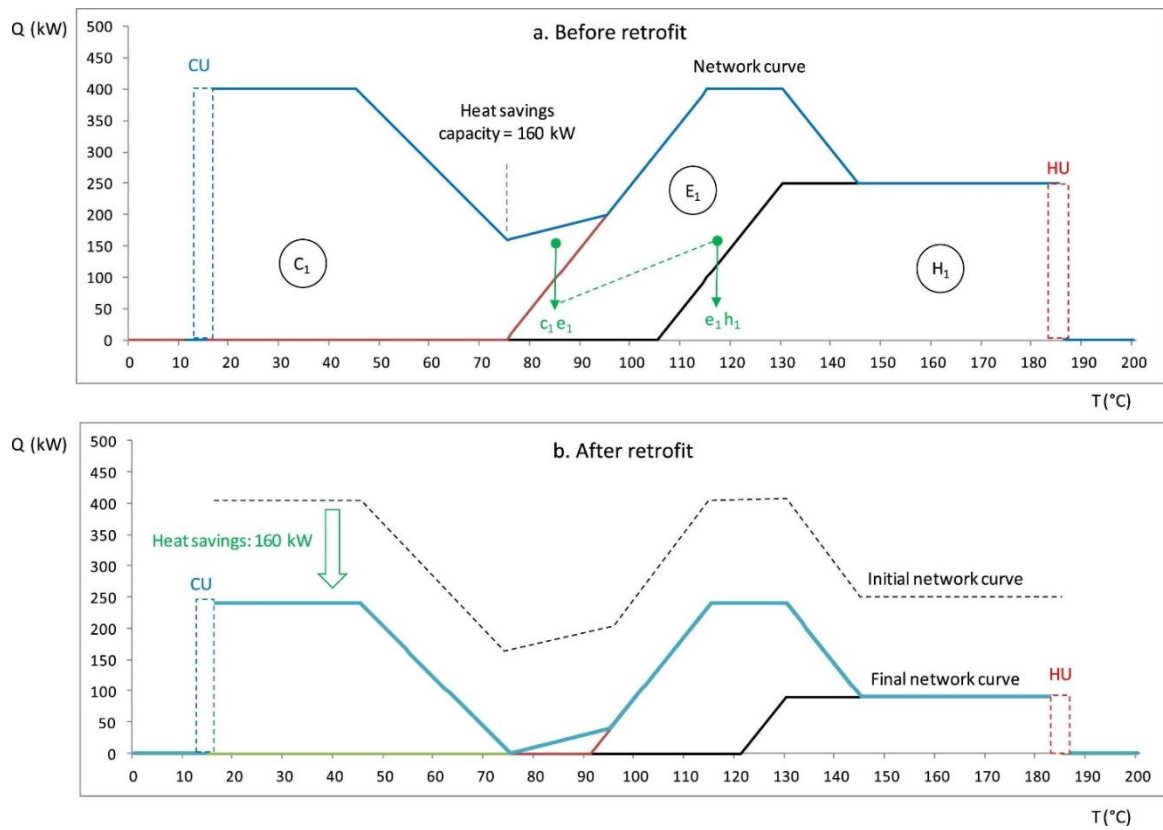


Figure 12 Original Energy Transfer Diagram identifying a heat bridge (Bonhivers et al., 2017).

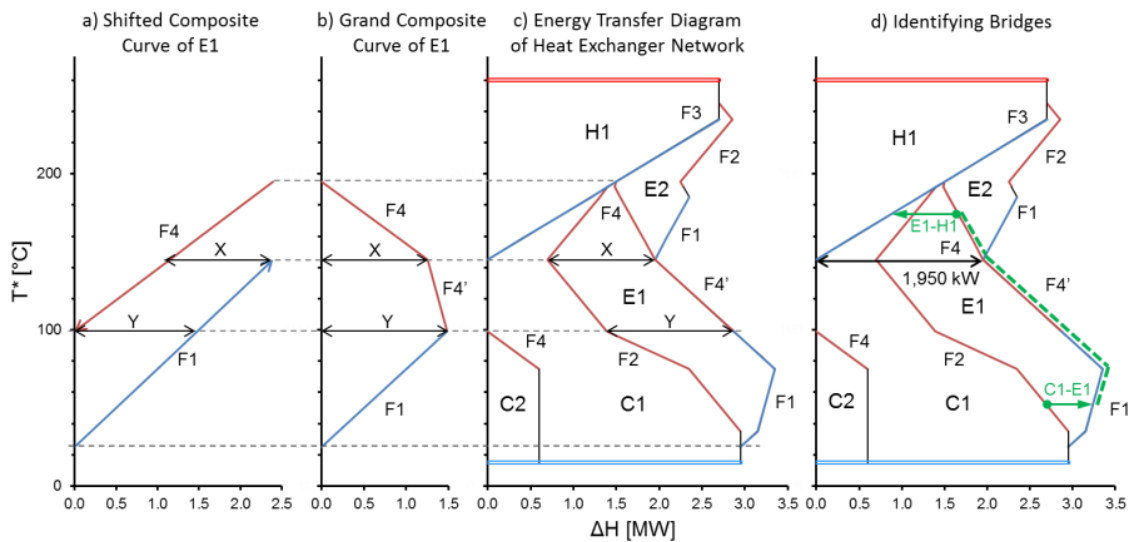


Figure 13 Developmental stages for creating a Modified Energy Transfer Diagram (M.R.W. Walmsley et al., 2017).

Heat Pump Bridge Analysis utilises the same methods and tools Modified Bridge Analysis (Schlosser et al., 2021) but also adds the capability in a retrofit bridge to match a heat pump operation to a heat surplus at a lower temperature interval with a heat deficit at a higher level at the cost of electrical work. The electrical work can be estimated using COP regression curves like those found in Schlosser et al (2020a). The paper also outlines three design guidelines: 1) the ΔT_{lift} of the heat pump should be minimised via smart HEN integration to increase the heat pump COP, 2) fully satisfy the demand of steam for individual streams to reduce the amount of utility HXs required. 3) Use the Pinch appropriate placement rules. The method was demonstrated into a case study, integrating a heat pump with a COP of 3.44 and removing all hot utility requirement and reducing required cooling utility from 850 kW to 320 kW (Figure 14).

(a)

T _i	T [*]	C2		C1		E1			
		25 kW/K	F4	15 kW/K	F2	25 kW/K	F4	20 kW/K	F1
	[°C]	[kW]		[kW]		[kW]		[kW]	
T ₁	122.5								
T ₂	117.5								
T ₃	97.5								
T ₄	95.85								
T ₅	92.5								
T ₆	72.5								
T ₇	49.5								
T ₈	37.5								
T ₉	17.5								
T ₁₀	12.5								
Σ		300		1175		1200		-1200	

(b)

T _i	T [*]	C2		C1		E1			
		25 kW/K	F4	15 kW/K	F2	25 kW/K	F4	20 kW/K	F1
	[°C]	[kW]		[kW]		[kW]		[kW]	
T ₁	122.5								
T ₂	117.5								
T ₃	99.15								
T ₄	97.5								
T ₅	95.85								
T ₆	92.5								
T ₇	85.5								
T ₈	72.5								
T ₉	49.5								
T ₁₀	37.5								
T ₁₁	28.75								
T ₁₂	17.5								
T ₁₃	12.5								
Σ		300		1175		1200		-1200	

Figure 14 The Heat Surplus deficit table being used to identify good bridge pathways (Schlosser et al., 2021)

2.4.3 Automated Heat and Work Retrofit Methods

One common issue the heat pump retrofit methods discussed thus far is that as the process increases in the number of steams, IB methods become cumbersome and at some point, unable to process the enormous number of possible available. This is where automated method becomes favourable. Most

methods are either MP based or a hybrid of the IB and MP that solve the network for an optimal solution based on a user set objective function. These methods have received praise and critique from the literature with the two most noted concerns being the lack of user input and insight (Tjoe & Linnhoff, 1986), and propensity to end up in local optima not finding the “best solution” (Klemeš & Kravanja, 2013). Currently, there are very few automated heat pump retrofit methods published in the literature.

Only two papers containing a dedicated heat pump retrofit method were found. Ahmetović et al. (2018) proposed an automated method for retrofit of MVR for evaporator systems. A general superstructure of an evaporator system was used that formulated a Mixed Integer Non-Linear Problem (MINLP). To solve this MINLP, a general algebraic modelling system was used to explore the different permutations. These permutations then become the optimisation problem to be solved in the next step of the method. The paper concludes by demonstrating that the method can integrate MVR and reduce the overall annualised cost for doing so, however, the method could not justify whether the solution was the global optimum and also took an immense amount of computational power.

Kang and Liu et al. (2015) used a multi-objective optimisation model to minimise total annual CO₂ emissions via retrofit of HEN and heat pump operations. It also added a pareto front analysis, acknowledging that an optimal solution can come in many forms and configurations, and therefore leaving it to a user to decide. Logical thermodynamic constraints were also included to limit the MP search space to only feasible solutions, saving on computational time. The limitation of this method was its ability to integrate different market available heat pumps, only considering bespoke simple single cycle, subcritical models.

2.5 Digitalisation and Digital Twins

The limitations identified in the preceding sections highlight the need for digitalisation within PI&E workflows. Process integration methods require accurate and structured representations of process streams, utility systems, heat exchanger networks, and thermophysical behaviour. Electrification further increases this requirement because heat pump performance depends on temperature lift, source and sink matching, and interactions between heat and work. A digital twin can therefore act as the connecting layer between process simulation and process integration, enabling stream data to be extracted consistently, process modifications to be evaluated iteratively, and electrification options to be assessed within the full process context.

The origins of the Digital Twin (DT) concept are often traced back to NASA's emergency response during the Apollo 13 mission in 1970 (Ferguson, 2020). Faced with a life-threatening systems failure, engineers relied on high-fidelity virtual simulations of the spacecraft to test possible interventions and identify a safe recovery strategy. Although the term "digital twin" had not yet been coined, this episode demonstrated the value of a digital counterpart to a physical asset for real-time decision support. The formal concept of the digital twin emerged much later, when (Grieves, 2014) described it as a system comprising a physical object, a digital representation, and a connection between the two domains. His later refinements distinguished between a digital twin prototype, instance, and aggregate, and NASA formally adopted the concept for space vehicle applications in 2010–2011.

Early academic papers of digital twins often equated them with high-fidelity simulations, with limited attention to the nature of the link between the physical and virtual domains. Subsequent work placed growing emphasis on the importance of dynamic, bi-directional connections (P. Wang et al., 2019), introducing notions of continuously updated replicas, and even self-adaptive or self-evolving systems (Tao et al., 2018). By the late 2010s, definitional refinements emerged to distinguish between levels of integration. (Kritzinger et al., 2018), for example, differentiated digital models (no automated data flow), digital shadows (one-way automated flow), and true digital twins (two-way automated flow). (Madni et al., 2019) extended this framing by classifying digital twins according to their fidelity, data integration, and use of machine learning for capturing operator behaviour or environmental dynamics.

A large proportion of published digital twin studies to date have focused on manufacturing systems, equipment monitoring, or individual devices, with comparatively few dedicated to process energy systems. Addressing this gap, Yu et al. (2022) provide one of the first systematic reviews of Energy Digital Twins (EDTs), offering a framework to classify and evaluate their role in industrial decarbonisation.

The review defines EDTs as a subset of digital twins aimed at minimising lifecycle energy use, enabling renewable integration, and reducing emissions, and introduces a multi-dimensional classification scheme. EDTs are categorised by their *likeness* (from simple block diagrams to full 3D models), *behaviour* (static vs. dynamic), *connectivity* (offline models vs. automated two-way twins), and *scale* (nano to macro). Analysis of more than 50 EDT studies shows that most remain at the digital model or shadow stage, with limited two-way integration with physical assets. Around 78% relied on simplified one-dimensional models, and the majority focused on process monitoring and optimisation rather than retrofit, lifecycle integration, or energy system transformation.

The authors also found that EDT applications are biased towards micro- and meso-scale units such as furnaces, reactors, or cooling towers, with little progress at the site or regional scale where renewable integration and energy symbiosis become most critical. They conclude that EDTs could be transformative for decarbonisation if developed with greater fidelity, lifecycle coverage, and scalability. To achieve this, they call for advances in (1) retrofit and service applications, (i2) macro-scale energy twins, (3) adaptive and self-evolving models, (4) secure data frameworks, and (5) AI- and hybrid-driven modelling approaches. Overall, Yu et al. (2022) position EDTs as a key enabler for net-zero energy systems, while emphasising that their application to process energy systems is still at an early stage

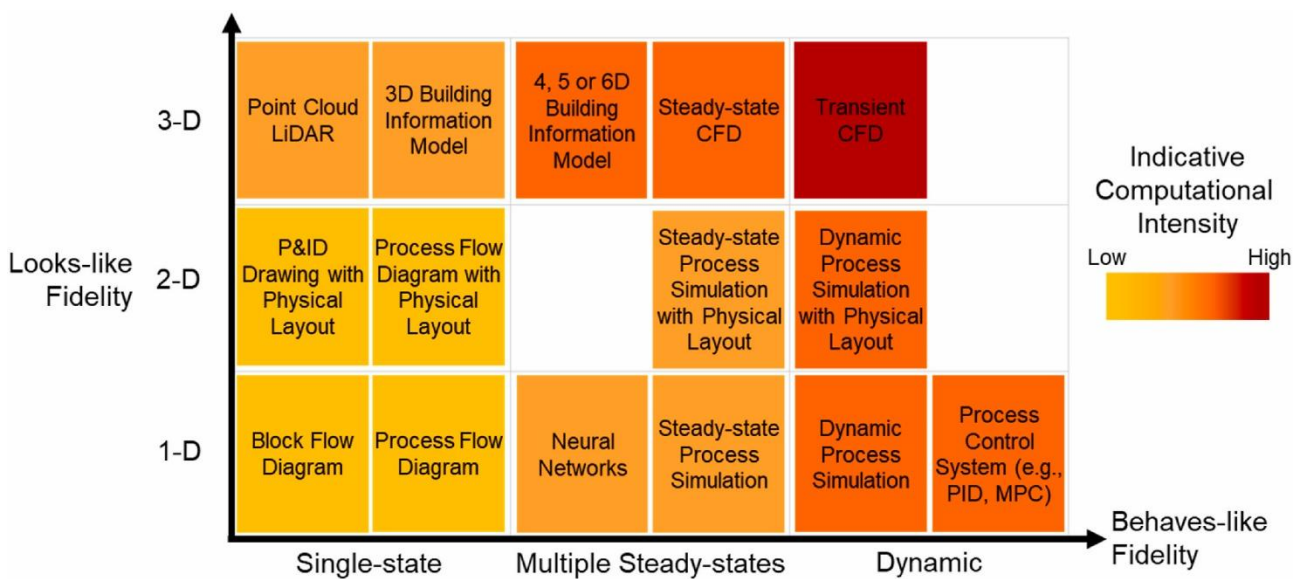


Figure 15 Classifying EDT families using the proposed classification framework figure from (Yu et al., 2022)

The recent systematic review by Ba et al. (2025) confirms that digital twins for energy efficiency remain a rapidly evolving research field with strong momentum across multiple industrial sectors. Following PRISMA guidelines, the authors screened over 500,000 publications, ultimately selecting 50 high-quality studies for detailed analysis. Their findings show that digital twin applications can deliver energy savings of up to 30%, reduce operational costs, and improve predictive maintenance, with particularly notable impacts in smart buildings, industrial manufacturing, and energy systems.

The review highlights that DTs are increasingly being integrated with Industry 4.0 technologies such as IoT, AI, and cloud computing, enabling real-time monitoring, scenario testing, and proactive optimisation of energy use. However, the field is still maturing major barriers include high implementation costs, interoperability with legacy systems, lack of standardised frameworks, and

cybersecurity concerns. The authors emphasise that more research is needed on scalable methodologies, cross-sector comparisons, and long-term performance validation. They argue that developing robust standards and integrating advanced analytics will be key to unlocking the full potential of energy digital twins.

Overall, this review demonstrates that energy digital twins are not only a hot research topic but also a strategic enabler for industrial decarbonisation and sustainability. Yet, the technology is still in a formative stage, with most case studies confined to early applications in buildings and smart grids. For process heat and energy-intensive industries (where emissions reductions are generally larger) applications remain sparse, highlighting a clear opportunity for further research and development.

2.5.1 A Standardised Energy Digital Twin Framework for Process Systems

A recurring theme in the recent mentioned reviews of energy digital twins is the absence of a standardised framework to guide their development and application. While numerous studies demonstrate potential energy and emissions benefits, the lack of common definitions, architectures, and protocols makes it difficult to replicate findings, and scale solutions across industries. This gap has been highlighted as a critical barrier to moving digital twins from conceptual pilots into reliable, industry-wide tools for decarbonisation.

E. Örs et al.(2020) propose a conceptual framework for Operational Digital Twins (ODTs) in the chemical process industry, with a strong emphasis on integrating artificial intelligence into process systems engineering. The framework is structured around five core building blocks: data management, process modelling, process optimisation, production scheduling, and process control, all of which are embedded in the advanced process control hierarchy. These elements are complemented by deployment and visualisation layers, which provide the interface between digital tools and real-time plant operation.

A key contribution of this framework is the explicit role assigned to AI methods. Surrogate modelling is used to approximate high-fidelity simulations, predictive modelling supports soft sensing and forecasting, and AI-assisted optimisation and control enable faster, more adaptive decision-making. This allows ODTs to deliver functions such as flexible operation, predictive maintenance, and multi-objective optimisation including CO₂ emissions reduction. The framework also points to the broader business implications of digitalisation, including the shift towards platform-as-a-service and software-as-a-service models for process industries.

While the framework provides a clear structure for linking process engineering with AI-based digitalisation, it remains largely conceptual, with limited demonstration through practical case studies. Issues such as uncertainty quantification, integration with modular plant design, and validation in retrofit-heavy industrial contexts are acknowledged as open research challenges. Nevertheless, this work represents an important step towards defining a structured methodology for ODTs in process industries, positioning AI as a central enabler of future industrial digital twins.

Pal et al. (2025) provided one of the most comprehensive reviews to date on digital twin (DT) applications within process industry, with a particular focus on frameworks and architectures. Their analysis identifies the essential components that underpin DT frameworks and links papers that give examples. The components and their corresponding articles that talk about them are: a physical layer of assets (reactors, distillation columns, heat exchangers) areas pioneered by companies such as Siemens, Ansys and Emerson (Qi et al., 2021) ; data collection and processing systems such as DCS/SCADA and IoT-enabled sensors; modelling and simulation layers combining first-principles and data-driven models example from (Rasheed et al., 2020); analytics and decision-making capabilities for predictive maintenance and optimisation example from (Sleiti et al., 2022); and visualisation and interface tools including dashboards and AR/VR environments for operator interaction. A defining feature across these frameworks is the presence of feedback loops that enable continuous synchronisation between physical and digital entities, ensuring real-time adaptability.

Building on the earlier reviews of digital twin frameworks, which highlight the persistent lack of a standardised, user-accessible frameworks, Burroughs et al. (2025) attempted to tackle this problem directly in their ESCAPE 35 contribution. The paper recognises that while advanced Process Systems Engineering (PSE) tools exist such as IDAES, Modelica, and GEKKO and that they offer powerful modelling and optimisation capabilities, their adoption is still constrained by steep learning curves, fragmented architectures, and poor accessibility for industrial users. These barriers limit the broader application of digitalisation and digital twins in energy-intensive sectors, particularly where decarbonisation and retrofit analysis require flexible, open, and integrative modelling environments.

To address this, the authors propose a new PSE+ ecosystem vision, combining the mathematical rigour of equation-based modelling with modern software engineering practices. A central principle of the ecosystem is user accessibility. The framework separates back-end modelling logic from the front-end interface, ensuring that users can interact through an intuitive UI rather than needing proficiency in programming languages or abstract algebraic modelling syntaxes. Behind the scenes, standardised data structures allow information from process models, thermodynamic property

packages, or external analyses (e.g., Pinch Analysis) to be stored consistently and reused across modules. This unified data layer enables different forms of analysis to operate on the same system description, promoting integration rather than fragmentation.

The framework also leverages cloud-native technologies such as containerisation and microservice orchestration. Each analysis module is deployed as a containerised service, ensuring compatibility across hardware and operating systems while enabling scalability through cloud infrastructure. Tools such as Kubernetes provide self-healing, load balancing, and parallelisation capabilities, meaning complex optimisation tasks can be distributed efficiently without compromising performance. This approach mirrors modern software practices in other digital industries but is rarely applied in process simulation environments.

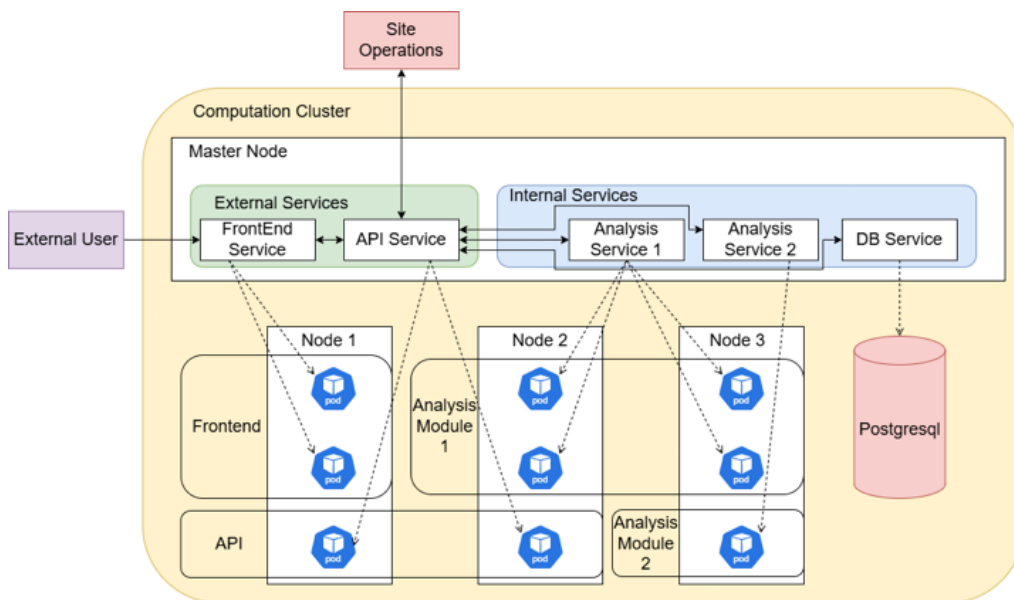


Figure 16 Microservice architecture deployed on Kubernetes sourced from (Burroughs et al., 2025)

Finally, the PSE+ ecosystem embeds extensibility and openness as design priorities. By building on open-source platforms (e.g., IDAES (Lee et al., 2021), Pyomo (Bynum et al., 2021)) and promoting modular contributions, the framework aims to avoid the vendor lock-in and rigid architectures of legacy simulators.

Although the functionality of the ADT incorporates many aspects, several limitations remain. Firstly, the framework is presented at a conceptual level. Whilst a substantial code base exists, it has had limited demonstration through industrial case studies. Without applied validation, it is unclear how well the proposed architecture can manage the complexities of real industrial data, particularly in retrofit-heavy sectors where process constraints, legacy equipment, and non-ideal fluid behaviour

present major challenges. Secondly, whilst modularity offers flexibility, it also risks creating fragmentation unless strong governance and standardisation can be established. Finally, the framework does not yet resolve one of the most persistent gaps in the literature: the development of specific property models and unit operations for non-petrochemical industries, such as and dairy, where complex fluids and semi-continuous processes are prevalent. Without addressing these modelling fundamentals, the platform risks reproducing the same limitations as commercial simulators, only in a more modernised software architecture.

2.5.2 Digital Twins for Milk Evaporation Plants

Milk evaporation plants are among the most energy-intensive subsystems in dairy processing (second behind the spray drier). Their complexity arises from the interplay of multiphase flows, vapour recompression technologies, and strong integration (recycling) between unit operations, making them a critical focus for electrification and process integration. Digital twins offer a powerful means of addressing this challenge, providing a virtual environment where evaporation processes can be modelled, optimised, and reconfigured to support both greenfield design and retrofit scenarios.

In terms of literature for “digital twins” for multi effect evaporation only one paper was found. Soares et al. (2019) developed and used a digital twin to address the persistent challenge of limited instrumentation and control in industrial multi-effect evaporation (MEE) systems. It was however formulated around the sugar industry rather than dairy. The authors postulated that conventional operation often relies on manual decision-making by supervisory staff, leading to suboptimal performance and inefficient energy use. The authors propose a digital twin framework that links a dynamic process model with live plant data to enable real-time monitoring, improved decision-making, and enhanced process stability.

The digital twin is built on a set of dynamic mass and energy balances, complemented by empirical equations for boiling point elevation, heat transfer coefficients, and thermodynamic properties of juice and steam. State observers were introduced to compensate for the lack of direct measurements, correcting flows and pressures using available setpoints and operator knowledge. This design allows the digital twin to function with minimal additional instrumentation, a crucial factor for deployment in existing industrial plants.

Validation was performed using industrial data from a four-effect evaporation train. The digital twin demonstrated strong agreement with measured values for pressure, levels, and particularly brix

concentrations, which are critical indicators of evaporation performance. Importantly, the system maintained numerical stability during transients, including start-ups and shutdowns, and was able to detect sensor miscalibration. These features position the tool not only as a monitoring system but also as a soft sensor and training environment.

Despite these contributions, limitations remain. The model was tuned manually without optimisation-based parameter estimation, which may limit predictive accuracy for different operating regimes. Furthermore, while effective for monitoring, the framework stops short of offering integration with optimisation or control functions that would enable closed-loop decision-making. Finally, although validated for sugarcane juice evaporation, the approach does not directly address the more complex thermophysical behaviours of dairy fluids, which represent a significant additional challenge for digital twin deployment in milk powder plants.

2.5.3 Review of Model-Based Optimization of Energy Consumption in Milk Evaporators

Although the term digital twin is now gaining traction, it is not widely used in the literature on evaporation systems. Instead, most studies in this area have been framed as process modelling and model-based optimisation exercises, in the context of improving the energy efficiency of milk evaporation.

Tsochatzidi et al. (2024) presents a comprehensive model-based study of falling film evaporators (FFEs) for milk concentration, focusing on strategies to reduce energy use and operating costs through thermal vapour recompression (TVR) and mechanical vapour recompression (MVR). The authors develop a dynamic mathematical model of FFEs, incorporating milk composition, thermophysical property correlations, mass and energy balances, heat transfer relations, and pressure drops. Five different evaporator layouts are analysed, ranging from pilot-scale two- and three-effect systems with TVR to industrial-scale designs and an MVR-based configuration.

The study employs global system analysis (GSA) with Monte Carlo uncertainty sampling to examine the influence of feed temperature, suction ratio, and discharge pressure on key process responses such as steam economy, product dry matter content, and annual cost. This is followed by dynamic optimisation under multiple scenarios to minimise steam cost, maximise yield, and reduce total annualised cost. Results demonstrate that increasing feed temperature and suction ratios can enhance steam economy and concentration, but at the expense of higher costs. For 50% solids (milk powder

production), three-effect TVR systems were most cost-effective, while for 30–35% solids (milk concentrates), MVR became favourable if electricity prices fell by ~11% or if gas-based steam prices rose by ~5%. Switching from powder to concentrate production yielded annual cost savings of 10–44%, while renewable-based steam (particularly biomass) showed potential for further cost reductions if market prices decrease.

While the study provides valuable insights into optimising existing TVR and MVR configurations, it does not propose any novel process designs that incorporate emerging electrification technologies such as heat pumps, which have been shown in other literature to offer significant potential for reducing both costs and emissions in evaporation systems. Instead, the analysis is limited to incremental improvements of conventional steam-driven configurations. Furthermore, the work does not explicitly address the retrofit context, even though most industrial decarbonisation efforts are applied to existing plants rather than greenfield facilities.

Kaviani et al. (2022) address the problem of high energy consumption in milk powder production, with particular focus on spray drying and associated thermal utilities, which are among the most energy-intensive operations in the dairy sector. The study applies pinch analysis using Aspen Energy Analyzer (AEA) to identify opportunities for heat integration and energy savings in a representative milk powder production cycle. After mapping hot and cold streams across the process (which included evaporators, spray dryers, cooling units, and heaters) the authors evaluated four heat exchanger network scenarios under constant heat recovery approach temperature assumptions.

The results demonstrate a pinch point of 85 °C, with minimum hot and cold utility requirements of 25.3 kW and 66.4 kW, respectively. Of the four network configurations tested, the MER-3 (Maximum Energy Recovery) option emerged as the most favourable, with five exchangers, a total heat transfer area of 65.3 m², and the lowest total annualised cost. This scenario achieved energy recovery close to theoretical targets while minimising investment costs, although at the expense of higher heat exchanger area compared to alternatives.

The paper's main contribution lies in demonstrating the value of systematic pinch-based optimisation for reducing dairy energy use and capital cost trade-offs in heat exchanger networks. However, several limitations remain. The approach is confined to conventional heat recovery between process streams and does not consider integration of emerging electrification technologies such as heat pumps or mechanical vapour recompression. Furthermore, the study is framed for greenfield design, with no

explicit consideration of retrofit challenges in existing plants where layout, operability, and legacy utilities often constrain implementation.

Taken together, the literature demonstrates steady progress in applying process modelling, optimisation, and more recently digitalisation to the energy-intensive operations of the dairy sector. Studies such as Bühler et al. (2019) and Tsochatzidi et al.(2024) highlight the potential of vapour recompression and energy minimisation approaches, while works like Soares et al. (2019) and Kaviani et al. (2022) show how modelling, monitoring, and Pinch Analysis can improve operational insight and identify cost-effective heat recovery strategies. More broadly, reviews of digital twin frameworks point to the transformative potential of digitalisation for decarbonisation. Yet across this body of work, important gaps remain: most studies emphasise incremental energy efficiency rather than novel electrification strategies such as heat pump integration; few explicitly address the retrofit context that dominates real industrial decarbonisation; and there is limited application of advanced targeting, advanced property modelling, and using the digital tools for rapid design iteration. These gaps establish the motivation for this thesis, which develops and demonstrates a platform-agnostic methodology that integrates exergy analysis, electrification strategies, and digital twin approaches to enable practical, retrofit-ready decarbonisation of milk evaporation systems.

2.6 Property Modelling for Complex Fluids

Thermophysical property models define the relationships between temperature, pressure, composition, and key thermodynamic quantities such as enthalpy, entropy, and density. These relationships underpin all mass and energy balances within the simulation and are therefore fundamental to every stage of PI&E analysis, including targeting, design and evaluation. Inaccuracies in property models can propagate through a generated digital twin, resulting in erroneous process targets, suboptimal technology selection, and unreliable techno-economic assessments. For electrification studies, where changes in efficiency or temperature lift can determine the economic viability of a retrofit or new design, property model precision is critical.

Commercial/conventional simulation environments, such as Aspen HYSYS, Aspen Plus (AspenTech, 2025), and DWSIM contain extensive libraries of pre-configured property packages. However, these models have been primarily developed for petrochemical, hydrocarbon, and high-temperature applications, where the fluids are comparatively simple and well-characterised. In contrast, food and dairy processing industries handle complex, multi-component, and often compositionally variable

fluids such as milk, whey, and humid air. These can exhibit strong non-idealities, phase transitions involving solids, boiling point elevation, and other complex behaviour.

A further limitation of commercial simulators is their restricted scope for property package customisation. While users can typically add new fluids, specify custom coefficients, or modify interaction parameters, introducing entirely new thermodynamic models is considerably more difficult due to the closed and opaque nature of proprietary architectures. Although some platforms allow the integration of user-developed models via dynamically linked libraries (DLLs), this approach requires specialised programming expertise, is dependent on the vendor's external model interface, and often lacks the flexibility to modify solver-level behaviour. Consequently, extending these tools to support advanced, solver-friendly models for novel or highly complex fluids remains a significant technical barrier, particularly for sectors such as food and dairy where existing property packages are fundamentally inadequate.

In open-source, equation-oriented modelling environments (such as the IDAES framework) these limitations are further compounded by specific numerical requirements of large-scale nonlinear optimisation (Lee et al., 2021). Solvers such as IPOPT (the default solver for IDAES) rely heavily on derivative information to navigate the solution space efficiently (Wächter & Biegler, 2006). The first derivative, or gradient, indicates the direction of steepest descent, guiding the search for improved solutions in gradient-based methods, such as IPOPT. The second derivative, or Hessian, captures curvature information, enabling methods to adjust both direction and step size to accelerate convergence, often achieving quadratic convergence near the solution. Without derivative information, algorithms must rely on less efficient approaches, often requiring many more iterations to achieve similar accuracy. As Wright (1999) notes, the availability of first and second derivatives transforms optimisation from a blind search into a mathematically guided process that leverages sensitivity information to achieve robust and rapid convergence.

A further complexity is that solvers typically assume the underlying functions are smooth and twice-continuously differentiable. This means both the first and second derivatives must vary continuously across the entire domain of interest. For many property functions especially those involving phase equilibrium, this requirement is challenging to satisfy. Phase transitions, by their nature, introduce discontinuities or non-smooth regions where phases appear or disappear depending on temperature, pressure, or composition. If these transitions are not handled carefully, numerical solvers can experience instability or fail to converge, particularly near phase boundaries where small changes in state conditions lead to large changes in phase composition or other properties.

The lack of suitable tools to generate these smooth, differentiable functions further compounds the difficulty. While options such as analytical differentiation, automatic differentiation, or numerical approximation of derivatives exist, each presents trade-offs between development effort, precision, and computational cost. Most existing property model libraries (particularly those embedded in commercial simulators) are not designed to meet these stringent requirements and therefore do not provide the level of derivative information needed for robust equation-oriented modelling. Platforms such as IDAES have therefore developed specialised property tools and methodologies to produce consistent and differentiable property functions. However, implementing these models demands significant investment in both development and validation, representing a substantial barrier for many users, especially those without deep expertise in thermodynamics and numerical methods (Lee et al., 2021)

2.6.1 Simulation of Milk for Evaporation

As described in the introduction design digital twins can be an indispensable tool for predicting process behaviour, guiding decision-making, and supporting optimisation without risking real plant operation. Commercial simulators such as Aspen Plus, HYSYS, VMGSim, and ProSimPlus are widely used in chemical and petrochemical industries because they provide extensive component libraries, robust thermodynamic packages, pre-configured unit operations, and validated computational frameworks (Munir et al., 2012). One of the key challenges in constructing these twins include the absence of dairy-specific components in standard libraries (Z. Wang & Hirai, 2011).

To address this, researchers have explored methods for representing milk as either a single pseudo-component or as a mixture of pseudo-components. Early work such as Bon et al. (2010) treated milk as a single component, which simplified simulation but limited applicability to processes such as pasteurisation or evaporation. More advanced studies such as Madoumier et al. (2015) modelled milk as a mixture of water, proteins, fats, lactose, and minerals, enabling the prediction of physical properties under different processing conditions. These approaches have gradually expanded the scope of simulation, from simple heating and cooling to more complex operations such as concentration and drying. Nevertheless, modelling challenges remain, particularly the prediction of viscosity at high total solids, the influence of whey proteins in heat treatment, and the need for accurate property correlations that capture both temperature and compositional effects.

Munier et al. (2016) attempted to address these limitations. A major strength of this work is its explicit treatment of both compositional and temperature effects on milk properties. As stated previous studies, either modelled milk as a single pseudo-component or omitted key constituents such as whey proteins, which limited applicability to processes like ultrafiltration or heat treatment. Munir et al. instead included casein and whey proteins, capturing their differing responses to heating, and employed expanded fluid and power-law models to address the non-Newtonian viscosity of concentrated milk. For heat capacity, the authors introduced a novel correlation capable of representing the double Gaussian peaks associated with fat melting transitions, allowing accurate simulation across fat contents up to 35%.

The validation against experimental and literature data showed high accuracy, with deviations of less than 2% for thermal conductivity, 4% for density, 3% for viscosity, and 1.5% for heat capacity. This level of fidelity demonstrated that pseudo-milk could be reliably embedded into commercial simulators (VMGSim in this study) to support mass and energy balance calculations in dairy processes such as evaporation and drying. Importantly, the method was flexible enough to handle different milk types (whole, skim, and concentrates), expanding the scope of process design and optimisation studies beyond what earlier models allowed.

The pseudo-milk framework set a foundation for integrating accurate thermophysical properties into process simulators, bridging a critical gap between food engineering needs and the capabilities of generic chemical process tools. Some limitations were present within the work. First, the paper does not provide a VLE formulation for milk. Without an activity-coefficient/equation of state basis that links composition to water activity, boiling-point elevation (BPE) and phase splits must be imposed via empirical curves rather than arising from the model itself. That limits rigorous flash calculations and constrains evaporator design and vapour-recompression studies, where BPE couples directly to temperature lift and work. Second, the implementation targets a sequential-modular simulator, not an equation-oriented environment. To support digital twins on these platforms (such as the Ahuora Platform), the pseudo-milk correlations must be recast as smooth, differentiable property functions with exposed partial derivatives (and ideally second derivatives) and integrated VLE/BPE relations.

2.6.2 History of accurate property packages for vapour compression cycles

Accurate modelling of thermodynamic properties is even essential for the design, simulation, and optimisation of processes involving fluids across a wide range of conditions, including gas, liquid,

two-phase, and supercritical states. In many industrial applications a such as vapour compression cycles, power generation, and energy systems accuracy across all phases and across all properties (entropy especially) be computed with a level of accuracy that approaches experimental uncertainty (Huber et al., 2022). This is particularly critical when small errors in enthalpy, entropy, or density predictions can lead to significant deviations in device performance, energy efficiency, or control outcomes.

Earlier modelling approaches for systems such as vapour compression cycles relied on simplified or semi-empirical equations of state (EoS), such as the Carnahan–Starling–DeSantis (CSD) (De Santis et al., 1976) or Extended Corresponding States (ECS) (Leland & Chappellear, 1968) models. These methods were advantageous when data availability was limited, as they required relatively few parameters and could represent the general behaviour of both vapour and liquid phases. However, they lacked the precision necessary for modern applications and often exhibited large deviations from experimental data, especially in the near-critical and superheated regions or for polar fluids to which vapour compression cycles operate commonly (Huber et al., 2022).

More advanced forms, such as the Modified Benedict–Webb–Rubin (MBWR) EoS (Jacobsen & Stewart, 1973), improved on this by incorporating a larger number of parameters and fitting to a broader set of thermodynamic data, including density and heat capacities. While MBWR models achieved a higher level of accuracy, they imposed limitations on the mathematical structure of the equations and required more extensive data sets for parameterisation.

To overcome these challenges, Helmholtz-based equations of state were introduced. These models define the Helmholtz free energy as a fundamental function of temperature and density, from which all other thermodynamic properties such as pressure, enthalpy, entropy, speed of sound and more can be rigorously derived using thermodynamic identities. The formulation also provided a further increase in accuracy for representing experimental data. This is exemplified in Figure 17 which shows a comparison of ECS, MBWR and Helmholtz EoSs. The figure shows the half black dots (representing Helmholtz) form much closer to the zero-deviation line than the other EoS models.

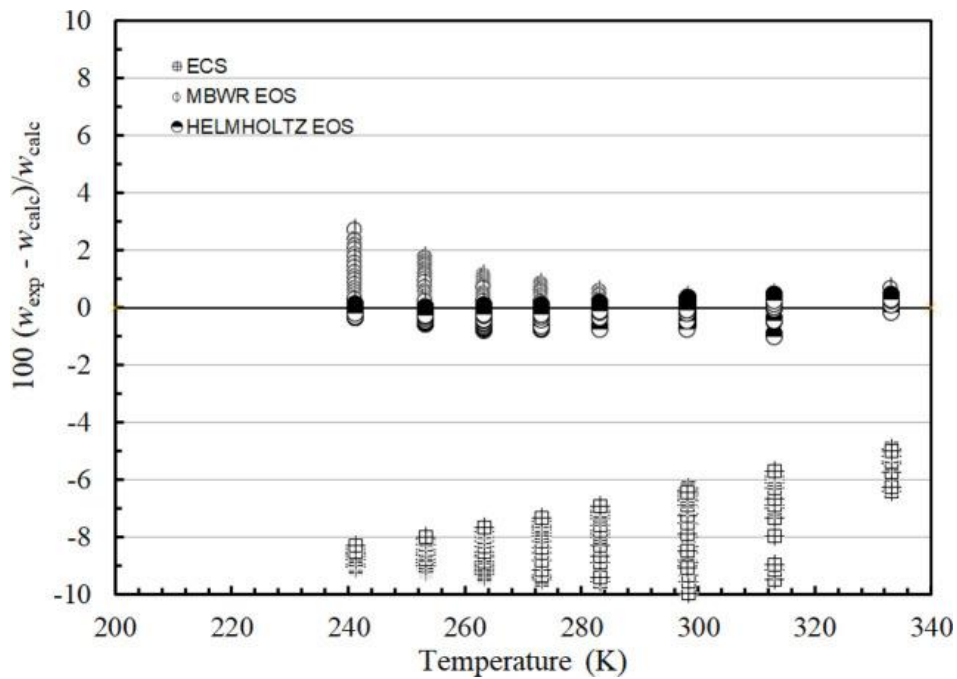


Figure 17 Deviations in the representation of the sound speed data of Takagi (Takagi, 1996) for R-125 (figure from (Huber et al., 2022))

2.6.3 Current tools and challenges for implementation

One implementation of the Helmholtz-based property models is found in the REFPROP software developed by the U.S. National Institute of Standards and Technology (NIST) (Huber et al., 2022). REFPROP employs multi-parameter Helmholtz equations to represent the properties of a wide range of industrial fluids, including refrigerants, hydrocarbons, cryogenics, and working fluids used in power and thermal systems. It supports extremely accurate calculations across all thermodynamic states and is used as a reference standard in many engineering applications. However, REFPROP is a proprietary tool, which can limit accessibility and extensibility for some users (Huber et al., 2022).

In contrast, CoolProp is an open-source thermophysical property library that also includes Helmholtz-based models for selected fluids (Bell et al., 2014). Originally developed to offer REFPROP-like capabilities in an accessible format, CoolProp supports integration with programming languages such as Python, C++, MATLAB, and Excel. Although it does not match REFPROP's coverage or accuracy for all fluids, it has become a popular alternative for academic and engineering applications requiring transparent, user-extensible models (Bell et al., 2014).

Commercial process simulators such as Aspen Plus and Aspen HYSYS are widely used in the chemical and energy industries; they primarily rely on cubic equations of state and activity coefficient

models. Some refrigerants within these platforms may use internally parameterised Helmholtz models, but users are generally not able to extend or modify them directly. Engineering Equation Solver (EES), commonly used for educational and HVAC applications, also employs Helmholtz models for certain fluids, although it functions as a closed-source desktop environment.

Despite the advantages offered by commercially available Helmholtz-based property models, challenges related to licensing restrictions, limited user expertise, and restricted accessibility persist for this category of high-accuracy formulations. One notable exception is CoolProp, an open source thermophysical property library that provides prepackaged, callable functions alongside transparent and modifiable source code. However, to fully integrate CoolProp into equation-oriented simulation frameworks such as the Ahuora software platform (Burroughs et al., 2025) it is necessary for the property model to expose, at a minimum, the first-order partial derivatives of its outputs with respect to the relevant state variables (Burgard et al., 2018). For example, the partial derivatives of enthalpy with respect to temperature and density must be available to ensure compatibility with the solver's optimisation routines. Without this functionality, the integration of CoolProp into such modelling environments is significantly constrained.

2.6.4 Implementation of Helmholtz into the IDAES equation orientated framework

Conveniently, the IDAES framework provides some tools to assist in the implementation of a Helmholtz equation of state. The key tool utilised is a library of prepackaged external function calls. This is important because Helmholtz equations of state themselves are high-order polynomials (seen later in the chapter) which require specialized root finding routines to give the right answer for phase equilibrium and variable state variable inputs. Implementing the property calculations as external functions allows the problem to be decomposed where specialized techniques can be used to ensure that the property models solve reliably and correctly, while reducing the complexity of the process model in which they are used. This library includes some common refrigerants fluids and the International Association of the Properties of Water and Steam (IAPWS-95) formulation of water (Wagner & Pruß, 2002).

Table 2 List of fluids translated into the equation-orientated Helmholtz formulations; uses and references

Fluid	Usage	Reference Data
H2O (Natural)	All-purpose fluid	(Wagner & Pruß, 2002)
CO2 (Natural)	Low GWP trans-critical high glide applications	(Span & Wagner, 1996)
Propane (Natural)	Low GWP HVAC, Refrig and HP	(Lemmon et al., 2009)
R1234ze(e)	Low GWP lower flammability HVAC, Refrig and HP	(Thol & Lemmon, 2016)
R125	Legacy high GWP for HVAC and flame suppression	(Lemmon & Jacobsen, 2005)
R134a	Automotive air-con very, high GWP, non-flammable	(Tillner-Roth & Baehr, 1994)
R227ea	High temperature applications, high GWP, non-flammable	(Lemmon & Span, 2015)
R32	Commercial and residential aircon, medium GWP	(Tillner-Roth & Yokozeki, 1997)

The availability of Helmholtz-based property models for a wide range of conventional refrigerants is highly beneficial for process and mechanical engineers seeking to simulate existing vapour compression systems. These fluids, many of which are included in legacy installations, continue to represent most currently deployed systems. However, with increasing regulatory and environmental pressures to decarbonise industrial sites and buildings, refrigerants with medium to high global warming potentials (GWPs) are being phased out. In response, the adoption of low-GWP alternatives such as ammonia, butane, and other natural refrigerants is being considered. To effectively apply the Process Integration and Electrification (PI&E) methodology to emerging refrigeration and heat pump configurations, it is essential that accurate property packages for these natural fluids are included.

2.6.5 State of the art modelling of Humid Air

The American Society of Heating, Refrigerating and Air-Conditioning Engineers (ASHRAE) has undertaken several major research projects to update psychrometric data and models. Historically,

the standard formulations in the ASHRAE Handbook Fundamentals have been based on data from Hyland & Wexler (1983). However, advances in measurement techniques and fundamental theory over the past two decades have prompted significant updates. These include the adoption of new equations for dry air (Lemmon et al., 2000), revised enhancement factors using updated Henry's constants (IAPWS, 2004), refined values for the molar gas constant (Mohr & Taylor, 2005), and improvements in water's equation of state and phase boundaries (Wagner & Pruß, 2002). Additional updates include enhanced air-water interaction models (Herrmann et al., 2009).

Due to the focus of this section being of surrogate modelling for complex systems only a brief explanation will be provided on how the humid air EoS is formulated to so the reader the complexity of the task. Further explanation is provided by Herrmann et al. (2009) for those who want a full explanation.

2.6.6 Formulation of the Humid Air Equation of State

Humid air can be modelled as a mixture of air and water vapor. In the simplest analysis, water and air are treated as ideal gases but in principle there is interaction between the air and water molecules that must be included using interaction parameters (similar to the modelling of milk previous).

The model is based on a virial-type EoS truncated at the third order, which enables the accurate representation of real-gas behaviour. The general form of the EoS is expressed as:

$$\frac{p\bar{v}}{RT} = 1 + \frac{B_m}{\bar{v}} + \frac{C_m}{\bar{v}^2}$$

where:

- p is the pressure,
- \bar{v} is the molar volume of the humid air mixture,
- T is temperature,
- R is the universal gas constant,
- B_m and C_m are the second and third virial coefficients for the moist air mixture.

The coefficients B_m and C_m are functions of temperature and composition and are calculated from the individual and cross-component interactions represented by the second and third virial coefficients (ψ_w is the molar fraction of water).

$$B_m = (1 - \psi_w)^2 B_{aa} + 2(1 - \psi_w)\psi_w B_{aw} + \psi_w^2 B_{ww}$$

and

$$C_m = (1 - \psi_w)^3 C_{aaa} + 3(1 - \psi_w)^2 \psi_w C_{aaw} + 3(1 - \psi_w)\psi_w^2 C_{aww} + \psi_w^3 C_{www}$$

For solving purposes, the equations derivative is also required and are described in the paper.

A key innovation in this model and complication is the comprehensive treatment of enhancement factors, which adjust the partial pressure of water vapor in the presence of air. The enhancement factor f , used in the expression for saturation pressure under mixed conditions, is computed from a detailed thermodynamic relationship involving virial coefficients, and Henry's constant from IAPWS (2004). This formulation corrects for deviations from ideal behaviour, particularly at elevated pressures and humidities which are of applications such as spray dryers.

It is only a function of temperature and pressure, but it must be iteratively obtained due to the nature of the expression for the enhancement factor causing issues for fully equation-oriented frameworks.

$$\begin{aligned}
\ln(f) = & \left[\frac{(1 + k_T p_{w,s})(p - p_{w,s}) - k_T \frac{(p^2 - p_{w,s}^2)}{2}}{RT} \right] \bar{v}_{w,s} + \ln[1 - \beta_H(1 - \psi_{w,s})p] \\
& + \left[\frac{(1 - \psi_{w,s})^2 p}{RT} B_{aa} \right. \\
& - 2 \left(\frac{(1 - \psi_{w,s})^2 p}{RT} \right) \left(B_{aw} - \left[\frac{(p - p_{w,s} - (1 - \psi_{w,s})^2 p)}{RT} \right] B_{ww} \right) \\
& + \left[\frac{(1 - \psi_{w,s})^3 p^2}{(RT)^2} C_{aaa} + \frac{3(1 - \psi_{w,s})^2 (1 - 2(1 - \psi_{w,s}))^2}{2(RT)^2} C_{aaw} \right. \\
& + \left. \frac{3(1 - \psi_{w,s})^2 \psi_{w,s}^2 p^2}{(RT)^2} C_{aww} + \frac{(3 - 2\psi_{w,s}) \psi_{w,s}^2 p^2}{2(RT)^2} C_{www} \right] \\
& + \left[\frac{(1 - \psi_{w,s})^2 (-2 + 3\psi_{w,s}) \psi_{w,s}^2 p^2}{(RT)^2} B_{aa} B_{ww} \right. \\
& + \frac{2(1 - \psi_{w,s})^3 (-1 + 3\psi_{w,s}) p^2}{(RT)^2} B_{aa} B_{aw} + \frac{6(1 - \psi_{w,s})^2 \psi_{w,s}^2 p^2}{(RT)^2} B_{ww} B_{aw} \\
& + \frac{3(1 - \psi_{w,s})^4 p^2}{2(RT)^2} B_{aa}^2 + \frac{2(1 - \psi_{w,s})^2 \psi_{w,s} (-2 + 3\psi_{w,s})^2 p^2}{(RT)^2} B_{aw}^2 \\
& \left. + \frac{p_{w,s}^2 (4 - 3\psi_{w,s})^2 \psi_{w,s}^3 p^2}{2(RT)^2} B_{ww}^2 \right]
\end{aligned}$$

Dry air properties are calculated using the fundamental Helmholtz energy equation from Lemmon et al. (2000) while the thermophysical properties of water (liquid and vapor) are sourced from IAPWS-95, and IAPWS-06 for ice. This multi-source integration ensures that each phase is modelled using the most accurate available formulation. Such properties would have to be implemented for use in the platform.

The resulting model is valid over a broad domain: temperatures from 130 K to 623.15 K, pressures from 0.01 kPa to 10 MPa, and humidity ratios up to 10 kg_w/kg_a.

Such a method would take a significantly long time to implement either by grey box modelling or by hand. The model is also discontinuous (non-smooth) and house complex formulations that requires

powerful iteration to solve. This makes Humid Air an ideal case study for the implementation of surrogate properties.

2.6.7 Machine learning for surrogate properties packages

Since the origins of the field, both data-driven methods and mathematical optimisation have been core pillars of PSE. More recently, advances in data science and machine learning have spurred a new wave of research that integrates statistical learning into mathematical optimisation frameworks (Ning & You, 2019), expanding the traditional boundaries of PSE. These developments have produced a spectrum of techniques, ranging from derivative-free optimisation (Rios & Sahinidis, 2013) to hybrid data-driven/mechanistic modelling (Boukouvala et al., 2016), and the use of surrogate models in optimisation (McBride & Sundmacher, 2019).

Misener & Biegler (2023) outlines three perspectives for developing process engineering surrogates. The surrogate-led perspective begins by choosing a surrogate model (e.g., polynomial regression, Gaussian processes) based on its statistical or smoothness properties and then formulates optimisation strategies tailored to that surrogate. The mathematical programming-led perspective instead prioritises optimisation properties such as linearity or convexity, selecting surrogates that align with the requirements of discrete or mixed-integer optimisation problems where global solutions are needed. Finally, the paper emphasises the verification challenge of ensuring that solutions derived from surrogate models are consistent with the underlying “truth” model and presents an optimisation strategy with demonstrated application to heat exchanger network synthesis.

The surrogate paradigms presented are representative of the surrogate models in literature of process system applications. However for the most part surrogate research has focused on unit and flowsheet models as exemplified in Ma et al. (2022) for extractive distillation and Goldstein et al. (2022) for amine scrubbing flowsheets. Within the context of Process Integration and Electrification (PI&E), machine learning provides a promising pathway for addressing one of the most persistent barriers: the lack of accurate, solver-compatible property packages for complex fluids.

An early surrogate model techniques for thermophysical properties was to simplify and abstract less significant variables to the thermodynamic model. Huber et al. (2008) utilised this approach by developed a surrogate mixture model to represent the thermophysical properties of a coal-derived aviation fuel (CDF) using a combination of gas chromatography–mass spectrometry (GC–MS) analysis and advanced distillation curve (ADC) metrology. The study recognised that CDFs, like

many alternative fuels, are chemically complex and difficult to characterise using conventional pseudo-component approaches common in petroleum modelling. Instead, the authors constructed a five-component surrogate mixture (n-propylcyclohexane, trans-decalin, α -methyldecalin, bicyclohexane, and n-hexadecane), selected to capture the volatility profile and chemical family distribution identified in the experimental analyses.

Each surrogate component was parameterised using Helmholtz-energy-based equations of state (Span–Wagner form) combined with viscosity models capable of representing full fluid surfaces rather than limited correlations. This allowed rigorous prediction of density, viscosity, and speed of sound. Validation against limited experimental data for the real coal-derived fuel demonstrated good fidelity: predicted deviations were within 1% for density, 4% for speed of sound, and 5% for viscosity. The authors also showed that pure trans-decalin alone could approximate many properties within ~1–8%, underscoring its dominance in the mixture.

Another method for surrogate modelling is a more data driven approach. Madana Gopal et al. (2023) review the challenges of thermophysical property modelling for cryogenic fluids under supercritical conditions, with emphasis on their importance in computational fluid dynamics (CFD) simulations of propulsion systems and advanced energy applications. Cryogenic fluids such as liquid nitrogen, oxygen, methane, and hydrogen exhibit highly non-linear variations in density, viscosity, thermal conductivity, and specific heat capacity near the critical point.

The paper first systematically evaluates the performance of real-fluid equations of state (EoS) (for example Peng Robinson and SRK) against NIST reference data. The authors conclude that SRK outperforms PR in the low-temperature liquid region, whereas PR is more accurate near the pseudo boiling transition.

Beyond EoS-based methods, the review discusses data-driven approaches, such as polynomial regression of NIST data and property tabulation, which improve computational efficiency in CFD while retaining acceptable accuracy. The authors did note fundamental limitations in applying common real-fluid EoS to fluids with strong quantum effects (e.g., hydrogen, helium), where density and transport property predictions diverge substantially from experimental data.

An additional and notable limitation is the reliance on polynomial regression models for property prediction, which, while computationally efficient, are poorly suited to representing highly non-ideal systems.

Amusat et al. (2024) present a novel equation-oriented (EO) modelling framework that integrates detailed water chemistry with process-scale cost optimisation for desalination systems. A central innovation of the work is the development and integration of multidimensional surrogate models to represent precipitation, pH, and mineral scaling behaviour within an optimisation environment. Instead of directly embedding complex electrolyte models (e.g., Pitzer, eNRTL, MSE) or relying on computationally expensive external calls to OLI Systems, the authors generate large datasets of simulated chemistry and fit surrogate models that are solver compatible. This allows mineral scaling constraints to be explicitly enforced in a rigorous optimisation framework while maintaining tractability.

The surrogates are primarily constructed using radial basis function (RBF) models with adaptive sampling strategies. For soda ash softening, simple one-dimensional surrogates accurately captured calcium carbonate precipitation and pH as a function of dosing, with $R^2 > 0.99$ and mean absolute errors below 1 mg/L. For the more nonlinear recarbonation and mineral scaling surrogates, adaptive sampling proved essential: the recarbonation model achieved $R^2 > 0.999$ with errors below 0.006, while the four-dimensional mineral scaling surrogates-maintained classification accuracies above 99% across thousands of scenarios. Importantly, the adaptive sampling scheme prioritised accuracy in the region near the scaling threshold ($ST = 1$), ensuring robustness of optimisation results where design decisions are most sensitive.

The study demonstrates how these surrogates can be seamlessly embedded in the open-source EO platform, enabling cost optimisation of high-recovery reverse osmosis systems under mineral scaling constraints. Results highlight how surrogate-enabled optimisation can uncover cost-optimal trade-offs between pretreatment (e.g., soda ash dosing, CO_2 recarbonation) and RO design, showing that at very high recoveries, pretreatment costs can exceed those of the desalination step itself.

However, limitations remain while the RBF approach works well here, it may face scalability challenges if extended to even higher-dimensional problems or when applied to complex process fluid models. Finally, as with most surrogate approaches, the fidelity of the optimisation is bounded by the quality and range of the training data.

2.7 Conclusions

This review has demonstrated that while the potential for process electrification is high, the practical challenges of integration remain substantial. Technologies such as heat pumps, mechanical vapour recompression, and electrode boilers are capable of supplying the majority of low- to medium-temperature process heat demands, yet their deployment is often constrained by temperature lift requirements, retrofit complexity, and economic trade-offs. Comprehensive studies of dairy processes remain rare, and those that do exist highlight both the technical feasibility of electrification and the significant barriers to its widespread adoption.

In terms of methodology, conventional process integration techniques, retrofit methods, and exergy analysis provide a valuable foundation but are insufficient in their current form. Classical pinch analysis offers useful energy targets but does not address the work requirements of electrification technologies. Retrofit strategies have advanced through tools such as modified bridge analysis, yet an inability to consider process equipment (like the evaporators) can lead to opportunities being missed. Exergy-based methods improve thermodynamic insight but require refinement and broader integration with simulation-based approaches. A further consideration is that the efficiency of electrification technologies, particularly heat pumps, is strongly dependent on how optimally they are integrated. Poor integration often results in excessive temperature lifts, which reduce the coefficient of performance (COP) and undermine both the thermodynamic and economic benefits of electrification. At the same time, iterative approaches that couple process integration with simulation or digital twin evaluation remain limited in practice. While such methods can, in principle, provide more rigorous assessments of operability and cost, their application has been narrow, often confined to case-specific demonstrations rather than generalised frameworks. As a result, there is a need for methodologies that combine the systematic targeting of integration methods with the rigorous, iterative validation capabilities of digital twins, ensuring that electrification strategies are not only theoretically efficient but also practical, operable, and cost effective.

The literature also highlights the limitations of both legacy simulators and equation-oriented frameworks when applied to food and dairy contexts. Commercial simulators have been developed primarily for petrochemical applications, and their property packages are poorly suited to complex fluids such as milk concentrates or humid air. Equation-oriented frameworks, while powerful for optimisation and digital twin development, demand differentiable and solver-compatible property

models that are rarely available for food systems. These constraints limit the ability of current tools to support electrification and process integration at an industrial scale.

A further limitation identified in the literature is the modelling of thermophysical properties for dairy fluids. Only a limited number of methods exist for representing milk and milk concentrates, and those that have been developed are often bespoke, non-standardised, and difficult to extend across different processes or compositions. Approaches such as pseudo-component models and surrogate-based formulations have shown promise, but they remain fragmented and are rarely implemented in widely used simulation environments. In addition, equation-oriented platforms introduce their own challenges, as they require property models that are smooth, differentiable, and solver-compatible. Many food-related property models fail to meet these criteria, creating barriers to embedding them in optimisation and digital twin workflows. To advance electrification and process integration in the food sector, methods must therefore be developed that both simplify the implementation of accurate property models and provide robust pathways for their integration into equation-oriented environments.

2.7.1 Gaps in the literature

From the review, several key gaps emerge that motivate the development of the methodology proposed in this thesis:

Incomplete integration of electrification into PI frameworks

While electrification technologies such as industrial heat pumps, MVR, and electrode boilers have been demonstrated in case studies, their integration into process integration frameworks remains limited. Existing methods typically prioritise energy minimisation but do not explicitly consider work requirements, temperature lift, or COP optimisation, which are central to the effective deployment of heat pumps.

Limited application of exergy-based targeting in practice

Exergy analysis provides a rigorous thermodynamic foundation for evaluating electrification strategies, but it is rarely embedded into systematic design or retrofit workflows.

Fragmentation in thermophysical property modelling for food and dairy systems

Legacy simulators lack accurate property packages for complex fluids such as milk concentrates and humid air. Existing surrogate or pseudo-component models are non-standardised, bespoke, and difficult to embed in equation-oriented platforms, limiting their use for optimisation and digital twin applications.

Limited deployment of digital twins for process integration and electrification

Although digital twins are recognised as transformative tools, most applications to date focus on monitoring, control, or asset-level optimisation in manufacturing. There is little evidence of comprehensive PI&E workflows being embedded within digital twin environments, particularly for food and dairy sectors.

Barriers to accessibility and adoption of advanced PI&E methods

Even where advanced methods exist, they are often locked in proprietary software, require significant thermodynamic expertise, or are demonstrated only in research contexts. There remains a need for open-source, modular digital tools that make PI&E methods more accessible to practitioners in non-petrochemical industries.

Chapter 3 Overarching Methodology

3.1 Introduction

The decarbonisation of industrial process heat represents one of the most pressing challenges in the global energy transition. In low- to medium-temperature sectors (such as food, beverage, and pulp and paper processing) thermal energy demand is substantial and is still overwhelmingly supplied by fossil fuels. This reliance persists despite the availability of electrification technologies, including industrial heat pumps, mechanical vapour recompression, and direct electric heating, which are capable of both improving energy efficiency and reducing emissions when effectively integrated. Achieving widespread adoption of these technologies requires systematic methodologies that can identify viable opportunities, quantify thermodynamic and economic trade-offs, and translate abstract targets into practical process designs.

This thesis responds to that need by developing a generalisable Process Integration and Electrification (PI&E) methodology, designed to support both greenfield design and retrofit applications. The methodology integrates exergy-based targeting, advanced heat and power integration strategies, and techno-economic assessment into a unified, stepwise framework. A central feature of the approach is its reliance on design digital twins (high-fidelity, flexible, and data-driven models that capture the complex interactions of industrial systems more accurately than traditional simulation tools). Digital twins provide the necessary foundation for rigorous targeting, technology screening, and iterative optimisation, ensuring that electrification strategies are not only thermodynamically favourable but also technically feasible and economically viable.

The chapter begins by presenting the overarching framework, which is structured into four iterative stages: (1) digitalisation of the process and development of a digital twin, (2) process integration and electrification targeting, (3) opportunity identification and concept development, and (4) detailed design evaluation. Together, these stages establish a pathway from problem definition to validated process solutions, underpinned by consistent data flows and feedback loops.

To demonstrate its applicability, the methodology is applied to a case study in the New Zealand dairy sector, focusing on the milk evaporation system as a representative example of a thermally intensive sub-process. This sector provides an ideal testbed due to its high dependence on fossil fuels, significant contribution to industrial emissions, and strong alignment with available electrification

technologies. While the case study is dairy-specific, the framework is designed to be broadly transferable to other low- to medium-temperature process industries.

In doing so, this chapter establishes the methodological foundation for the remainder of the thesis, connecting the wider goals of industrial decarbonisation with the practical steps required to implement electrification at process level.

3.2 Methodology

The methodology developed in this thesis provides a structured framework for accelerating the decarbonisation of low- to medium-temperature process industries through Process Integration and Electrification (PI&E) supported by design digital twins. It addresses the dual challenge of improving thermal efficiency while enabling the substitution of fossil-based utilities with electrified alternatives such as mechanical vapour recompression (MVR), industrial heat pumps, and direct electric heating.

At its core, the methodology integrates three elements that have historically been applied in isolation:

1. Digitalisation – the creation of high-fidelity, flexible digital twins that capture process and utility behaviour at a level suitable for electrification studies.
2. Process Integration – the systematic identification of thermal recovery and utility reduction opportunities, using both energy and exergy-based targeting methods.
3. Electrification Design and Evaluation – the screening, development, and detailed assessment of electrification concepts to ensure thermodynamic feasibility, operational practicality, and economic viability.

The overall framework is structured in two parts:

- Part A: Creating a Digital Twin – This part focuses on creating a digital twin of the process. It covers preparation and scope definition (Chapter 4), the development of advanced thermophysical property packages for complex fluids (Chapter 5), and the construction of unit operations and system-level models to form a robust design digital twin (Chapter 6).
- Part B: Process Integration and Electrification – This part applies the digital twin as a foundation for PI&E. It includes exergy-based targeting, identification of electrification

opportunities, conceptual and detailed integration design, and techno-economic evaluation in both greenfield (Chapter 7) and retrofit (Chapter 8) contexts.

Together, these two parts form a platform-agnostic methodology that is generalisable beyond the case study presented in this thesis (dairy evaporation) and applicable to a wide range of low- to medium-temperature process industries.

3.3 Part A: Digital Twin Development

3.3.1 Purpose

The first stage of the methodology is the development of a digital twin of the process. A digital twin is a validated computational model that replicates the physical and thermodynamic behaviour of an industrial system. In contrast to conventional process simulations, which are often limited to steady-state representations, a digital twin is constructed to reflect the interactions between process units, utility systems, and operating conditions in sufficient detail to support analysis of electrification options.

The purpose of this stage is to establish a reliable and flexible modelling environment that can be used as the basis for subsequent integration and electrification studies. This requires not only a representation of process flows but also the capacity to account for the complexities characteristic of low- to medium-temperature industries, including multicomponent fluids, variable operation, and strong coupling between thermal systems. By providing this capability, the digital twin forms the technical foundation for the methodology.

3.3.2 Approach

The digital twin is developed in three steps, described in Chapters 4 to 6.

The first step involves preparation and scope definition. This includes identifying the objectives of the model, setting system boundaries, and collecting relevant process and utility data. Data are reconciled where inconsistencies occur, and the system definition ensures that the model reflects the physical process and its operating constraints. At this stage, attention is also given to identifying parameters that are expected to influence performance most strongly, so that the model can later be applied with confidence to integration and electrification analyses.

The second step concerns the development of thermophysical property models. Accurate property data are essential for ensuring that the model correctly represents mass and energy balances as well as phase and transport behaviour. For many industrial fluids, particularly those outside the petrochemical sector, suitable property models are not readily available, and their development can be challenging. In such cases, it may be necessary to combine fundamental thermodynamic approaches with modern digital tools. Data-driven methods, including machine learning, can support this process by extending predictive capability into regions where experimental data are sparse or highly variable. The objective is to obtain property models that are sufficiently accurate, robust, and computationally tractable to support the broader aims of the digital twin.

The final step is the construction of the system model. Unit operations are represented in detail and integrated into a flowsheet that describes the process as a whole. The model is validated against industrial data and benchmark cases to confirm that it reproduces observed behaviour. Where appropriate, the digital twin is configured to represent steady-state, dynamic, or multi-state operation, reflecting the requirements of the process under study. The completed model then provides a consistent platform for subsequent process integration and electrification analyses.

3.3.3 Conclusion

The outcome of Part A is a validated digital twin of the case study process, developed through structured data preparation, property model formulation, and system modelling. The digital twin enables accurate representation of energy and exergy flows, provides a consistent data pipeline framework for assessing integration and electrification opportunities, and establishes a workflow that is applicable beyond the specific case studied in this thesis.

In summary, the digital twin is not an end in itself but a prerequisite for the analyses that follow. It ensures that the subsequent stages of the methodology are based on a reliable representation of the process, thereby supporting systematic evaluation of integration and electrification strategies through easily accessible data.

3.4 Part B: Process Integration and Electrification Method

3.4.1 Purpose

Part B of the methodology applies the digital twin to the systematic development and assessment of electrification opportunities. The purpose of this stage is to determine how industrial processes can be reconfigured to minimise fossil fuel use and improve efficiency while ensuring that designs remain technically feasible and economically viable. This requires establishing thermodynamic performance targets, identifying integration pathways, and evaluating candidate technologies in a structured and iterative manner. The figure presented in this section illustrates the logic of the methodology, showing how process optimisation, targeting, and electrification design are combined into a coherent and iterative workflow.

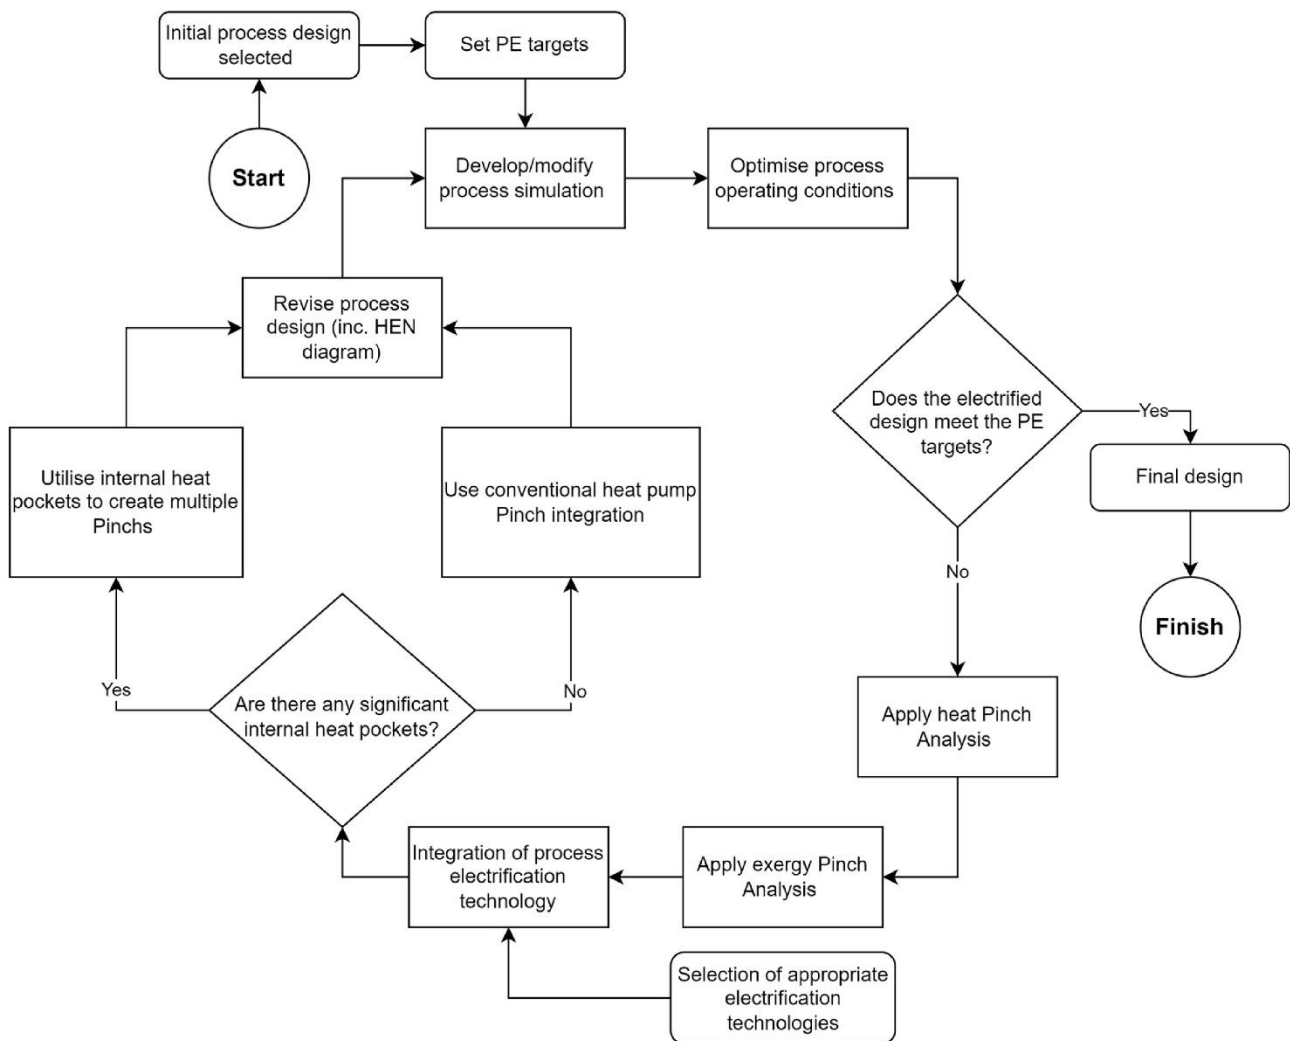


Figure 18 Process Integration and Electrification (PI&E) method.

3.4.2 Approach

The application of the methodology begins with the use of the digital twin to explore the operational performance of the process. Before new technologies are introduced, the existing system is examined to determine whether improvements can be achieved by adjusting operating parameters such as temperatures, pressures, flow rates, and recirculation ratios. This optimisation stage is essential, as it ensures that the baseline process operates as efficiently as possible and provides a clear reference against which electrification options can be evaluated. By systematically adjusting operating conditions within the constraints of product quality, process safety, and equipment limitations, it is possible to reduce utility consumption and identify new integration opportunities that might not be visible in the original configuration.

If, after optimisation, the process design still does not meet performance expectations for energy efficiency or electrification readiness, further analysis is carried out using integration techniques. Heat pinch analysis is first applied to establish the minimum heating and cooling demands of the system and to identify opportunities for internal energy recovery. This analysis provides a structured basis for revising the process design, including adjustments to the heat exchanger network, to align more closely with the thermodynamic limits.

Where electrification technologies are under consideration, exergy-based analysis is employed to provide a more comprehensive view of energy quality and work requirements. Exergy pinch analysis enables the identification of regions in the process where exergy destruction occurs and highlights where technologies such as mechanical vapour recompression or high-temperature heat pumps may be most effectively integrated. This step extends conventional targeting by ensuring that integration strategies consider not only the quantity of heat available but also its thermodynamic potential.

Based on the outcomes of the targeting and optimisation stages, candidate electrification technologies are selected for integration. Each technology is evaluated for technical compatibility with the process conditions, alignment with the identified heat recovery opportunities, and potential to meet performance benchmarks such as efficiency, emissions reduction, and electricity demand. In some processes, the presence of significant internal heat pockets may permit the creation of multiple pinch points, leading to more complex but potentially more efficient integration strategies. In other cases, simpler single-pinch configurations may be appropriate.

Once candidate technologies are identified, they are incorporated into the digital twin to develop revised process designs. The integration of these technologies requires detailed modifications to the system configuration, including updates to the heat exchanger network and adjustments to utility systems. The performance of each revised design is then evaluated within the digital twin, with results compared against the targets defined for energy, exergy, and electrification. The methodology is iterative in nature: if a revised design fails to meet these targets, the process configuration is revisited, and further modifications are explored until a satisfactory design is achieved.

3.4.3 Subtleties of retrofit

Electrification of existing industrial plants through retrofit introduces a distinct set of challenges compared to greenfield design. While the methodology outlined in Part B applies to both contexts, retrofit studies require additional steps and considerations to account for the constraints of existing equipment, utility systems, and site integration. The flowchart in Figure 19 illustrates the adapted workflow for retrofit analysis, which incorporates these additional layers of complexity.

The first subtlety arises from the need to align new electrification strategies with the existing process configuration. In retrofit scenarios, the system cannot be redesigned from first principles; instead, changes must be introduced incrementally while maintaining operability and compliance with production requirements. This often limits the scope of modification and places a premium on identifying opportunities for improved utilisation of waste heat within the existing heat exchanger network (HEN). Accordingly, the analysis begins with the construction of a comprehensive process heat balance and representation of the system using the heat surplus–deficit table (HSDT). This ensures that the retrofit opportunities are identified within the thermodynamic and physical constraints of the existing process.

A second important consideration is the role of waste heat recovery. In retrofit applications, the availability of low-grade heat is often significant, but its temperature may not be sufficient to directly displace fossil-fuel utilities. The methodology therefore incorporates a structured search for ways to unlock higher-temperature heat through modifications to the HEN. By applying retrofit strategies such as exchanger re-piping, stream splitting, or network reconfiguration, it is possible to elevate the utility value of waste heat and expand the integration potential for electrification technologies.

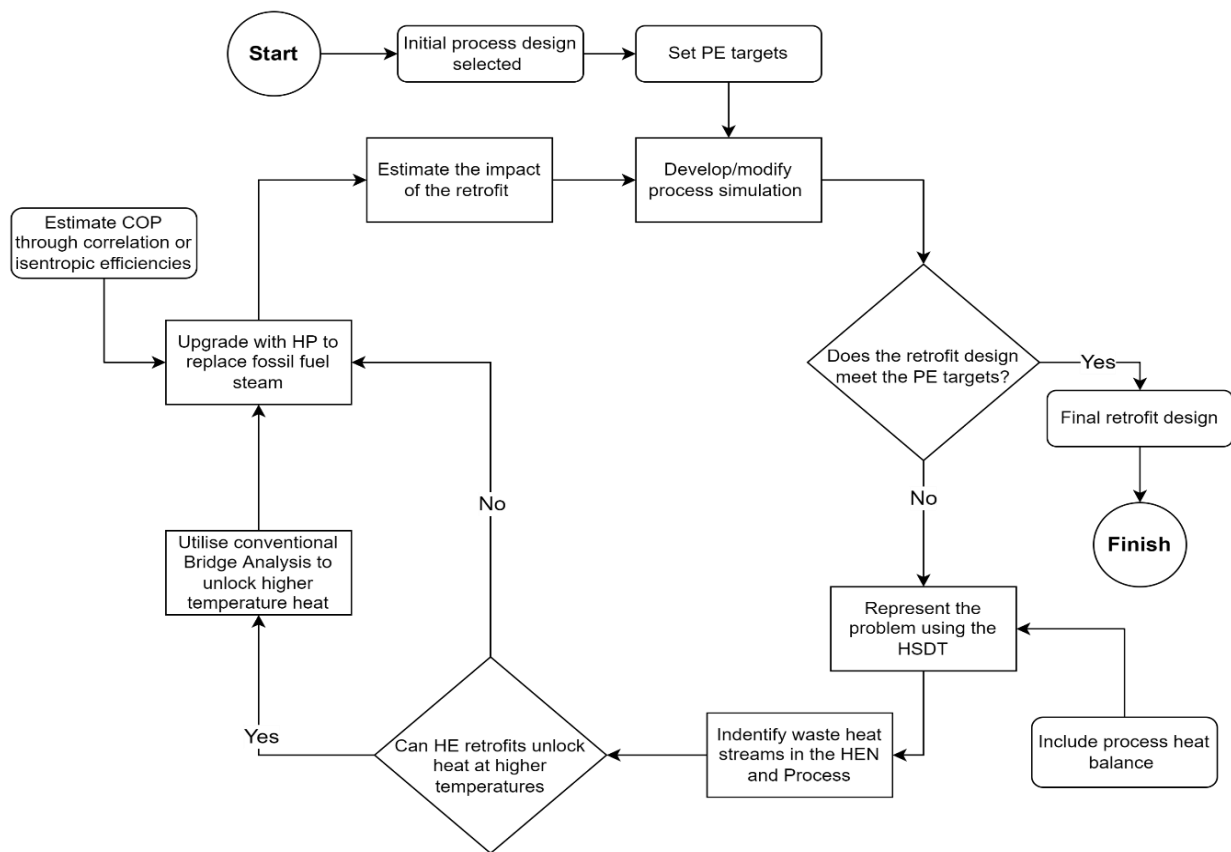


Figure 19 Process Integration and Electrification retrofit design method through incremental changes.

If heat exchanger retrofits cannot provide sufficient high-temperature recovery, electrification technologies such as industrial heat pumps are evaluated as an alternative pathway. Again, rather than assuming idealised performance, the coefficient of performance (COP) of retrofit heat pumps is estimated using either empirical correlations or isentropic efficiency models, reflecting the practical limitations imposed by equipment specification and integration constraints. This ensures that the electrification concepts remain grounded in realistic technical performance rather than optimistic assumptions.

Finally, and again, retrofit design must be assessed not only against thermodynamic targets but also in terms of its economic viability. A levelised cost of heat (LCOH) assessment is therefore embedded into the workflow to evaluate the cost-effectiveness of each retrofit option relative to fossil-based utilities. This is particularly important in retrofit contexts, where capital investment is constrained by the existing site and financial justification must be made against the incremental benefit of the retrofit rather than the total performance of a new system.

3.4.4 Conclusion

Part B has outlined how the digital twin is applied to process integration and electrification studies. In the case of greenfield design, the methodology provides a structured pathway from operational optimisation through thermodynamic targeting and technology integration to the development of new process configurations. In retrofit applications, the same framework is adapted to account for the constraints of existing equipment and utility systems, with additional emphasis on waste heat utilisation, realistic technology performance, and incremental economic justification.

3.5 Conclusions

This chapter has introduced the overarching methodology developed in this thesis for the process integration and electrification of low- to medium-temperature industries. The methodology is structured in two complementary parts. Part A described the creation of a digital twin as the essential foundation for systematic analysis. The digital twin provides a validated and flexible representation of the process, enabling accurate quantification of energy and exergy flows and supporting both steady-state and variable operating conditions. Part B outlined how the digital twin is then applied to integration and electrification studies. In this stage, operational optimisation, energy and exergy targeting, and systematic technology integration are combined within an iterative workflow to identify and refine electrification opportunities.

A key feature of the methodology is its applicability to both greenfield and retrofit contexts. In greenfield design, the framework provides a structured route from problem definition to new process configurations. In retrofit applications, the same structure is adapted to the constraints of existing plants, with particular attention to the reconfiguration of heat exchanger networks, the unlocking of higher-value waste heat, and the realistic evaluation of technology performance and economic viability. The result is a methodology that is generalisable across industrial settings while retaining the flexibility to address the specific requirements of each application.

The next chapter, Chapter 4, begins Part A of the methodology in detail. It focuses on the preparation of the digital twin, including the definition of system boundaries, data collection, and scope clarification. This stage establishes the foundation upon which the subsequent property model development (Chapter 5) and system modelling (Chapter 6) are built. Together, these chapters complete the digitalisation of the process and provide the platform required for the integration and electrification studies that follow.

Part A: Digital Twin Development

In this first part of the thesis, the focus shifts from conceptual methodology to implementation of the first phase: the development of a digital twin. Specifically, it describes how a design digital twin of a milk evaporation system was created to support electrification and energy integration studies. Unlike conventional simulation models, which are often static, opaque, and limited in scope, the digital twin developed in this research captures the detailed thermodynamic, operational, and integration characteristics required to support practical electrification design and retrofit decisions.

The part begins by outlining the rationale and requirements for digitalisation in the context of PI&E, highlighting the limitations of traditional simulation environments and the specific needs of non-petrochemical industries such as dairy processing. It then describes the stepwise process used to develop the digital twin, including model scope definition, data gathering, thermophysical property package development, and unit operation modelling. A particular emphasis is placed on the technical challenges and solutions associated with modelling complex fluids (such as milk and humid air) in Chapter 5.

Finally, it explores the role of Digital Twins in enabling effective energy integration, using milk evaporation as a case study. It critically evaluates the limitations of existing simulation approaches and demonstrates how digital tools can bridge the gap between process modelling and integration analysis. In doing so, it positions digitalisation not merely as a technological advancement, but as a necessary enabler for the widespread and effective application of process integration in industrial decarbonisation efforts.

Through this case study, the part of the thesis demonstrates how Digital Twins can overcome key barriers in electrification design and enables accurate, iterative evaluation of energy and exergy performance. The digital twin not only supports process-level design but also provides a foundation for subsequent analysis stages, including targeting and integrated process electrification design (Chapter 6). In doing so, it highlights the critical role of digital tools in advancing the practical application of PI&E in real industrial contexts.

Chapter 4 Digital Twin Goals and Preparation

4.1 Introduction

Process integration remains one of the most effective strategies for improving thermal efficiency and reducing energy-related emissions within industrial systems. Despite decades of methodological development and demonstrated success in academic settings, its uptake in real-world engineering practice has often been constrained by a persistent challenge: the absence of a direct, structured connection between conventional process simulation outputs, the data required for integration analyses, along with a feedback loop from process integration back to the simulation.

In standard process simulation environments, engineers routinely construct detailed steady-state mass and energy balance models to evaluate process performance, maintain product specifications, and support operational decision-making. However, when applying integration techniques such as Pinch Analysis, exergy-based targeting, or utility system optimisation, these simulation outputs often fall short. The core issue lies in the way traditional simulation tools manage thermodynamic information, focusing on the behaviour of discrete unit operations, typically through proprietary or opaque property models, and aggregating thermal data in ways that obscure its utility for system-level energy analysis.

While flowrates, temperatures, and heat duties may be accurately reported, the underlying thermal data such as streamwise enthalpy changes, temperature–enthalpy profiles, and minimum temperature approaches, are often embedded within unit-level black boxes or exported in unstructured formats. This fragmentation complicates the extraction and transformation of data for integration studies, rendering the process manual, time-consuming, and error prone. Engineers are often required to reconstruct thermal models from scratch, identifying hot and cold streams, reconciling inconsistencies, and manually generating new datasets purely for integration purposes. These challenges are exacerbated in complex systems like dairy evaporation, where working fluids such as milk do not conform to standard property packages, and where advanced operations like direct steam injection (DSI) and flash evaporation introduce additional thermodynamic complexity.

In this context, the development of digital twins offers an alternative, inter-operable analysis tools. Unlike traditional simulation, digitalisation is not limited to the behavioural definition of process models; it encompasses the creation of integrated, query-able data architectures that support traceability, modularity, and cross-disciplinary decision-making (Burroughs et al., 2025). Digital

systems treat mass and energy balances not as closed simulation artefacts, but as connected data sources that feed directly into heat recovery targeting, utility optimisation, and retrofit design workflows.

A clear distinction must therefore be made between process simulation and Digital Twins. The former is typically static and confined to a specific software environment, solving predefined equations to predict process behaviour under assumed conditions. The latter represents a system-level transformation that links simulation models, historical plant data, property packages, and optimisation tools within a unified, interoperable platform.

For milk evaporation systems, a design Digital Twin is especially essential. These systems are tightly integrated, characterised by multiple effects, flash vessels, vapour recompression units, and significant process variability. Accurate analysis requires not only sophisticated models but also multi-steady-state adaptability to changes in milk composition, flow rate, and energy demand. By embedding custom thermophysical property packages for milk and related fluids within a Digital Twin, it becomes possible to simulate and evaluate system behaviour with a level of precision and transparency that traditional tools cannot offer.

4.2 Method

4.2.1 Preparing to build a digital twin

In the development of a digital twin for the design of a modern milk evaporator, thorough preparation is fundamental to ensuring that the resulting model is technically rigorous, fit for its intended purpose, and capable of supporting subsequent tasks such as process integration, optimisation, and techno-economic evaluation. This preparation phase involves the clear definition of design objectives, the systematic collection and reconciliation of input data, the selection of an appropriate simulation environment, and the establishment of a model structure that balances detail, against the user's time to create.

4.2.2 Defining design requirements and modelling boundaries

The initial task in preparing the simulation is to define its design intent. This includes specifying whether the model is intended for greenfield plant design, retrofit evaluation, energy targeting, control strategy development, or process optimisation.

In the case of designing a modern fully electric milk evaporator, the focus is on greenfield design with an emphasis on delivering a system that meets production and food safety requirements while only using a sufficiently minimal amount of electric utility. With this purpose in mind, the boundaries of the process must be carefully established. It is important for the digital twin to capture all essential stages, and therefore it is important for the process and its scope to be fully understood.

4.2.3 Operational requirements and design ideals

In addition to defining system boundaries, the preparation phase must also address the operational requirements and constraints that govern process design. These constraints typically reflect safety, product quality, or regulatory considerations and directly influence the selection of technologies, heat exchange strategies, and integration opportunities. For example, restrictions may prohibit direct contact between certain fluids, necessitating the use of intermediate heat transfer loops or specific equipment configurations. Similarly, some processes require rapid heating or cooling rates to protect product integrity, which imposes requirements on the type of unit operations included in the model and on the accuracy with which their thermal behaviour is represented. Other constraints may specify allowable residence times, pressure limits, or environmental conditions, all of which shape the feasible design space.

Alongside these hard constraints, a number of design ideals guide the preparation of the simulation model. A common objective is the minimisation of electricity and utility consumption while remaining within practical and economic limits. Achieving this often requires detailed modelling of advanced energy recovery technologies, such as recompression systems, heat pumps, or optimised utility networks, so that the model can credibly evaluate performance and identify opportunities for improvement. Another important consideration is process simplification, aimed at reducing capital investment, operating costs, and the physical footprint of the plant.

In addition, the localisation of energy integration (such as clustering heat recovery loops or electrification technologies) can enhance operability and control. This approach encourages a modular model structure in which sub-systems can be isolated, analysed, and optimised independently, while still reflecting their integration within the overall process. Finally, the simulation must provide sufficient flexibility to investigate variations in key operating parameters, since changes in pressure, temperature, or utility availability can significantly influence system performance and energy efficiency.

4.2.4 *Data collection*

Unlike retrofit studies, where operational measurements can be collected and reconciled against plant records, the development of a digital twin for a greenfield facility requires a different approach. In the absence of site-specific operating data, the modeller must construct a representative dataset that is both technically credible and aligned with the intended design objectives.

The process of generating such a dataset typically involves three steps:

1. Identify appropriate benchmarks: Published case studies, design guidelines, or industrial best-practice reports provide initial estimates for stream data, operating conditions, and utility demands. These sources offer reference points that capture the ranges of temperatures, pressures, compositions, and loadings encountered in real systems.
2. Normalise and structure the data: Because greenfield plants can vary in scale, benchmark data are typically normalised to a production basis (for example, per unit of throughput). Stream flow rates can then be scaled to the intended design capacity, while operating temperatures and pressures are generally treated as scale independent. At this stage, the data should be organised into a structured format that explicitly lists all process streams, their thermodynamic properties, and any utility or waste flows.
3. Establish representative operating conditions: Sensitivity analysis is used to explore how variations in key design parameters affect energy use, process feasibility, or economic performance. By identifying which parameters exert the greatest influence, a representative set of conditions can be selected that balances realism with robustness. Parameters with minor influence can be fixed at central or typical values, while those with major influence should be chosen conservatively to ensure the model remains valid under expected variability.

4.2.5 *Data for sensitivity analysis*

Process designs developed through Digital Twin simulations must demonstrate robustness across a range of potential operating conditions. Industrial systems rarely operate at a single fixed set of temperatures, pressures, or product specifications; instead, variability in raw material properties, equipment performance, and operational targets means that a credible design should remain feasible and efficient under different scenarios.

To assess this robustness, a parametric sensitivity analysis can be used on each or the final process design. A set of representative design variables was selected based on their expected influence on energy demand and utility consumption. Examples of such variables include:

- Operating temperatures within major unit operations, which affect heat transfer requirements and utility balances.
- Treatment or reaction set-points, often constrained by safety or product quality specifications.
- Final product specifications (e.g., concentration or purity), which govern the overall separation or conversion duty.
- Feedstock properties or pre-treatment conditions, which define the baseline load imposed on the process.

Each variable is to be varied independently across a defined range informed by literature values, design heuristics, or industrial data.

4.2.6 Selecting the appropriate simulation platform

The selection of the simulation platform is another crucial aspect of preparation, as it directly impacts the fidelity, flexibility, and usability of the resulting digital twin. In this research, both a commercial simulator (Aspen HYSYS) and open-sourced platforms (DWSim and IDAES, which is the engine that sits behind the Ahuora Digital Twin platform) were utilised each offering distinct advantages and limitations depending on the task.

Commercial simulators such as Aspen Plus, Aspen HYSYS, and SuperPro Designer are widely used across industry due to their mature interfaces, extensive component databases, and comprehensive libraries of standard unit operations. These tools are often favoured in industrial settings because they require relatively little programming knowledge, support rapid model development, and provide user-friendly environments for visualising and adjusting process flowsheets. For standard (typically petrochemical) process fluids and configurations, commercial tools can offer fast and reliable simulation results.

However, the strength of commercial simulators in conventional applications often becomes a limitation when more advanced or non-standard functionality is required. Their proprietary architectures can restrict access to internal model structures, making it difficult to implement custom

unit operations, novel thermodynamic formulations, or non-standard integration routines. Property models are often optimised for petrochemical fluids and lack flexibility for user-defined behaviour, especially in sectors like food and dairy where complex multicomponent fluids such as milk or whey are involved. Moreover, integration with external optimisation tools or open-sourced packages such as Pysmo (a surrogate modelling framework from IDAES (Lee et al., 2021)) and OpenPinch (an automated targeting software for heat and work (Walmsley, 2025)) can be cumbersome or reliant on costly additional modules.

In contrast, open-source, equation-oriented platforms such as IDAES, Modelica, or the Ahuora platform offer much greater flexibility for advanced process systems engineering tasks. These platforms allow users to define and manipulate unit operations at a granular level, implement custom thermophysical property packages, and integrate with optimisation solvers and external databases. IDAES, in particular, was designed for advanced optimisation problems and includes a number of open-sourced, free to use packages. It is also written in python and therefore can integrate with the large number of other python-based tools simply.

However, these benefits come with trade-offs. Open-source platforms typically require significantly more setup effort, greater mathematical and programming proficiency, and a deeper understanding of process modelling principles. Constructing a complete process model in IDAES, for example, requires defining state variables, writing configuration blocks, and ensuring thermodynamic consistency across unit operations, all of which can be time-consuming and error-prone for users unfamiliar with equation-oriented approaches.

In this research, both types of platforms were employed to demonstrate that the PI&E method is platform agnostic and can be applied across a wide range of available simulation environments. However, certain aspects of the methodology become more challenging to implement depending on the nature and limitations of the selected platform. These challenges, along with the strategies used to address them, are documented throughout the thesis.

4.3 Case study and data analysis

4.3.1 Design intent

The first outcome of the preparation stage was a clear articulation of the design intent for the digital twin. The case study is centred on a greenfield, fully electric milk evaporation facility designed to

process raw milk into concentrate suitable for downstream spray drying. The overarching intent is to develop a system that meets food safety and product quality requirements while operating with the entirety on low emissions electricity for external utility.

Crucially, this intent serves not only as a starting point but also as a reference framework throughout model development and analysis. Each modelling decision (whether in defining system boundaries, selecting property packages, or assessing utility configurations) must be evaluated against the original design objectives to avoid outcomes that are either unnecessarily complex and detached from practical feasibility, or overly simplified in ways that obscure critical integration opportunities. By explicitly returning to the stated design intent at each stage, the digital twin is kept both technically rigorous and aligned with the practical requirements of an electrified dairy evaporation system.

4.3.2 The milk evaporation process and design scope

The case study considered in this thesis is a modern milk evaporation system designed to concentrate raw milk prior to downstream drying. At a high level, the process consists of milk preheating, thermal treatment, multi-effect evaporation, vapour recompression, and final concentrate heating. The system also includes supporting utility interactions such as condensate recovery, CIP water heating, refrigeration, and cooling. Figure 21 illustrates the base-case process flow configuration adopted as the basis for the digital twin.

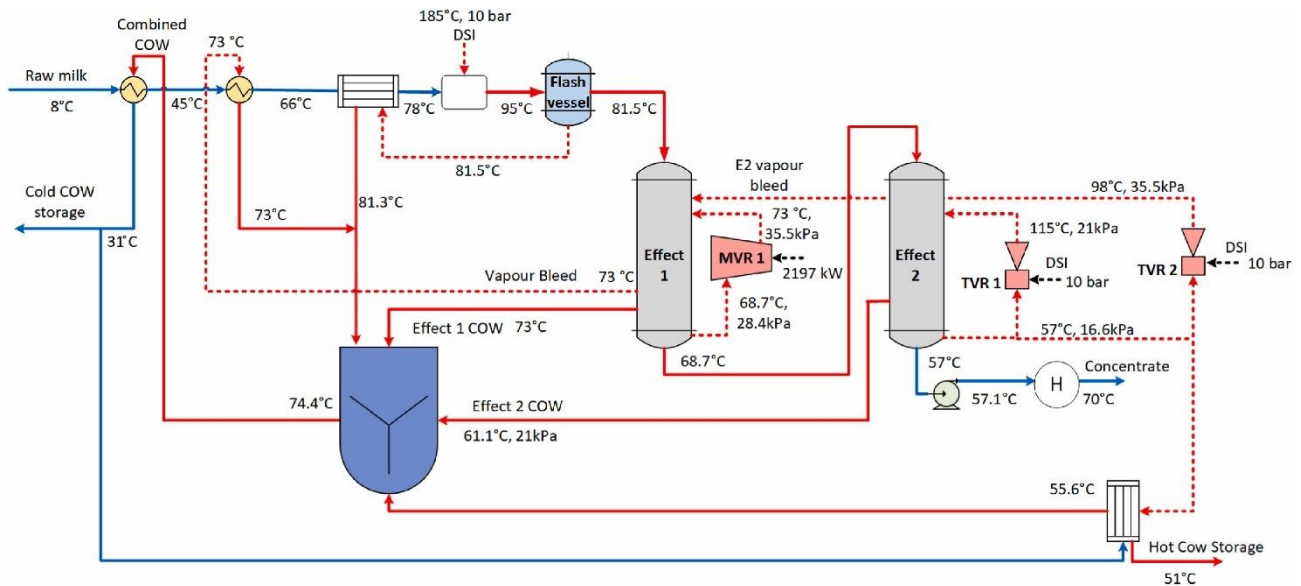


Figure 20 Base-case industrial milk evaporator integration illustrated using a process flow diagram

Within the evaporation train, the first effect is integrated with mechanical vapour recompression (MVR), while the second effect uses thermal vapour recompression (TVR). These features provide the reference structure for later electrification and integration studies. For the purposes of this chapter, the emphasis is on defining the process scope and model boundary rather than detailing the internal behaviour of each unit operation. Detailed modelling methods for individual components are presented later in Chapter 6.

The digital twin boundary therefore includes the main milk heating and evaporation stages, vapour recompression systems, concentrate heating, ancillary hot-water and cooling duties, and the relevant utility and waste streams associated with process operation.

4.3.3 Operational requirements and design outcomes

The preparation stage established a series of operational requirements that directly informed the configuration of the Digital Twin. A primary outcome was the definition of three critical constraints. First, segregation between process and utility fluids was required, necessitating the inclusion of intermediate heat exchange loops to eliminate any risk of cross-contamination. Second, high heating rates in the thermal treatment stage were shown to be essential, which in turn required the explicit representation of rapid heating technologies within the model. Third, the passage of product streams through the 45–60 °C range had to occur within a single exchanger to reduce the potential for microbial growth. These constraints collectively set clear boundaries on the choice of unit operations, the design of the heat exchanger network, and the permissible strategies for thermal integration.

In addition to these requirements, the preparation phase also defined a set of guiding design ideals. The most significant was the minimisation of electricity demand, achieved through the detailed modelling of vapour recompression, heat pump configurations, and utility systems. Process simplification was also prioritised, both to limit capital expenditure and to reduce plant footprint. Furthermore, the localisation of heat integration and electrification subsystems was emphasised to improve operability, allowing sub-systems to be optimised independently while still reflecting their integration within the wider process.

4.3.4 Data collection for the base case (Greenfield)

For a greenfield facility, no site-specific measurements exist to reconcile against. Instead, the preparation stage required the identification of a representative benchmark system that could provide realistic operating conditions, stream data, and process structure. To this end, the study adopted the case data reported by Walmsley et al (2016) in their work on appropriate placement of heat pump technologies in dairy evaporation systems. This reference provides a comprehensive set of thermodynamic and process data for a modern multi-effect milk evaporator, including vapour recompression units, utility requirements, and key process streams.

Their case fell within the scope outlined in the previous section. The site has a 2-effect evaporator system with integrated vapour recompression (Figure 20). The first effect is integrated with MVR and the second effect is integrated with TVR. The MVR effect currently operates at 68 °C (milk-side) and the TVR effect operates at 57 °C (milk-side). Both effects have approach temperatures of 5 °C (condensation temperature to evaporation temperature). Some vapour from the TVR effect is upgraded in a secondary thermocompression to indirectly pre-heat milk via the shell-side of the MVR effect.

Table 3 Base case site stream data.

Process stream	TS [°C]	TT [°C]	CP [kW/°C]	ΔH [kW]
Cold milk	8	95	272	23,557
HT flash vapour	83	83		3231
E1 vapour bleed	73	73		5733
E2 Condenser	56	56		1677
COW	73	13	237	14,103
CIP Water	15	85	9	644
Concentrate	56	70	49	692

Stream flow rates are determined by the scale of production, expressed as tonnes of powder produced per hour (tp/h). In contrast, most stream temperatures (including those defined by evaporator pressures and associated boiling points) are relatively independent of production scale and can be considered “soft” specifications. Certain temperatures, however, are fixed design parameters. These include the inlet temperature of standardised milk (8 °C), the heat treatment temperature (94.5 °C), the final outlet temperature of the milk concentrate (70 °C), and the supply and target temperatures

of CIP water (15 °C and 85 °C, respectively). These temperatures are industry standards and are vital for food safe product manufacturing.

Utility costs were based on New Zealand pricing, with steam and hot water heating costed at \$25.00 NZD per tonne and electricity at \$120 NZD per MWh. The complete set of site utility stream data is provided in Table 4. Heat transfer coefficients selected where: 2–8 kW m⁻² °C⁻¹ for liquid and vapour streams and approximately 0.1 kW m⁻² °C⁻¹ for gas streams.

Table 4 Base case site utility stream data

Utility stream	Type	TS [°C]	TT [°C]
HPS	Hot	250	250
MPS	Hot	180	180
HTHW	Hot	80	65
LTHW	Hot	55	35
Cooling Water	Cold	25	30

4.3.5 Data collection for sensitivity analysis (Greenfield)

Industrial-scale evaporators do not operate at identical conditions; operating temperatures, concentrations, and set-points can vary between plants depending on design choices and production requirements. Four key design variables were identified based on their strong influence on thermal performance and energy demand. Each variable was perturbed across a representative range derived from industrial practice and industrial data. The ranges examined are shown in Table 5.

Table 5 Sensitivity analysis of key design variables

Design variable	Range examined
Evaporator tube-side temperature	67–72 °C
Heat treatment temperature	90–100 °C
Final concentration	52–63 %
Initial concentration (after reverse osmosis pre-concentration)	14.5–30.5 %

This approach provides two key benefits. First, it identifies which design choices exert the greatest leverage on overall energy performance, thereby guiding priorities for optimisation. Second, it establishes bounds on the robustness of the proposed design by revealing conditions under which energy consumption may rise sharply. The analysis therefore ensures that the process model is not

only optimised for a single design point but remains technically and economically credible across a range of realistic operating conditions.

4.3.6 Selection of the simulation environment

Three platforms were evaluated for this thesis: the commercial sequential-modular simulator Aspen HYSYS, the open-source sequential-modular simulator DWSIM, and the Ahuora platform powered by the equation-oriented IDAES framework. Each environment offered distinct strengths and limitations that directly influenced the fidelity, flexibility, and usability of the resulting model.

Aspen HYSYS (AspenTech, 2025) provided the most rapid route to a functioning flowsheet. Its mature user interface, extensive libraries of standard unit operations, and integrated property methods allowed the evaporator process to be assembled and converged with relatively little development effort. These advantages explain its widespread use in industry. However, this ease of use came at a significant cost. Aspen HYSYS is a high-licence-fee commercial package, often prohibitive for smaller industrial users. In addition, its proprietary architecture limited transparency, and non-standard operations required workaround solutions implemented through custom spreadsheets or approximate correlations.

By contrast, DWSIM offered an accessible open-source environment that supported a wide range of conventional process modelling tasks. Its integration with CAPE-OPEN unit operations and availability of additional thermodynamic models made it attractive as a freely available alternative to commercial simulators. Nevertheless, there are limitations in model robustness and solver stability when applied to large-scale, recycle-intensive systems such as evaporators. In addition, the level of community support and validation for highly customised property methods was limited, reducing its suitability for advanced research applications.

The most comprehensive flexibility was obtained using the Ahuora platform powered by IDAES. This equation-oriented framework allows for the explicit definition of state variables, custom unit operations, and user-specified property packages. Integration with Python enabled connection to external optimisation routines, surrogate modelling tools, and process integration frameworks. The trade-off, however, was the substantially higher development effort required to construct and initialise flowsheets. Building a fully functional model demanded detailed specification of constraints, scaling factors, and solver configurations, making the environment less accessible to users without strong programming and mathematical modelling expertise.

Overall, the results demonstrated that no single environment was universally superior. Aspen HYSYS provided rapid development and industrial familiarity, DWSIM offered accessibility for conventional systems, and the Ahuora/IDAES platform delivered maximum flexibility and integration capability at the expense of effort and complexity.

4.4 Conclusions

This chapter has demonstrated that the successful development of a digital twin requires more than the assembly of a process flowsheet. Careful preparation is essential, beginning with a clear articulation of design intent that frames all subsequent modelling decisions. For the present case, the intent of creating a fully electrified evaporation system provided both a benchmark and a constraint, ensuring that the model remained aligned with the dual goals of product quality and low-emissions operation.

The definition of process scope and boundaries established the essential unit operations and ancillary systems that must be represented, while the identification of operational requirements and design ideals provided clear rules for modelling decisions. Hard constraints, such as fluid segregation, heating-rate requirements, and food safety considerations, set the limits of feasible design space, while guiding principles such as energy minimisation, process simplification, and modularity informed the selection of system configurations. Together, these outcomes ensured that the digital twin remained both technically rigorous and industrially credible.

In the absence of site-specific data, the development of a greenfield model necessitated the use of benchmark datasets and structured methods to generate representative operating conditions. The sensitivity analysis of key design parameters demonstrated that the system's energy performance is strongly influenced by a small number of critical variables, highlighting where future optimisation and electrification strategies should focus. This analysis also established robustness, showing that the model is capable of capturing a realistic range of industrial operating conditions rather than a single design point.

Finally, the selection of a simulation environment underscored the importance of matching platform capabilities to research objectives. Aspen HYSYS offered speed and industrial comparability but at high licensing cost and with limited flexibility. DWSIM provided accessibility and transparency but lacked robustness for complex recycle-intensive systems. The Ahuora/IDAES platform delivered the greatest flexibility and integration potential, albeit with a significant increase in development effort.

The thesis will utilise all three of the platforms to show the strengths and weaknesses of the different frameworks in Chapter 6.

In summary, the preparatory work described in this chapter established a robust foundation for the digital twin, ensuring that it is grounded in realistic operational constraints, informed by representative data, and supported by an appropriate modelling environment. These outcomes collectively enable the analyses presented in the subsequent chapters, where the digital twin is applied to explore process integration and electrification strategies in detail.

Chapter 5 Advanced Property Packages for Complex Fluids

5.1 Introduction

Accurate thermophysical property models are foundational to process simulation, but their importance becomes particularly pronounced in the development of digital twins for Process Integration and Electrification (PI&E). Property models define the enthalpies, entropies, phase equilibria, and transport properties that underpin each mass, energy and momentum balance. Errors in these models propagate through the simulation, compromising heat and material balances and, by extension, the validity of data passed on to energy targeting, optimisation, and system design. As stressed in Chapter 4, a digital twin requires an appropriate simulation environment supported by plant measurements. However, the accuracy and credibility of a design digital twin ultimately depend on the quality of the thermophysical property data embedded within it.

In conventional process simulators such as Aspen HYSYS, property packages are precompiled and optimised for common petrochemical fluids. While these tools are accessible and convenient for standard applications, most of their built-in models are rarely suitable for the complex, non-ideal, and compositionally variable fluids found in food and dairy processing. Moreover, they offer limited flexibility for implementing advanced user-defined formulations.

In contrast, platforms such as IDAES offer the flexibility to define entirely custom property models and access internal thermodynamic expressions. This flexibility introduces substantial complexity. IDAES, for example, is equation-orientated and requires the user to provide first and second derivatives of all property functions to ensure numerical stability and convergence within nonlinear optimisation solvers. In addition, solvers assume that property functions are continuously differentiable which is a requirement that is difficult to meet near phase transitions or in models involving discontinuous behaviour.

Developing thermodynamic property packages that meet these requirements poses a significant barrier to the practical implementation of digital twins, particularly in non-petrochemical industries where suitable models are not readily available. For users without deep expertise in thermodynamics and numerical methods, this challenge can stall adoption of even the most advanced simulation platforms.

In the milk evaporator case study, milk is the fluid with limited property packages available in open literature and commercial software. Raw bovine milk is a complex multi-component fluid, typically composed of approximately 87 wt% water and 13 wt% total solids. The solid fraction consists of several major constituents: fats (approximately 4 wt%), proteins (3.4 wt%), lactose (4.8 wt%), and mineral salts (0.8 wt%) (Zhang et al., 2014). These components exist in a finely emulsified and colloidal state, and their thermophysical properties vary significantly with temperature, concentration, and processing conditions.

This chapter details the development and implementation of advanced property packages required to create an accurate digital twin for a dairy powder plant. It also explores the use of machine learning techniques to support the integration of complex thermodynamic models within equation-oriented simulation environments. An outline of the chapter is included below.

1. **Creating Activity Coefficient model** for the simulation of liquid milk and its vapour.
2. **Custom Helmholtz-Based Formulations** for high-priority pure substances such as refrigerants and steam.
3. **Machine-Learned Surrogate Models** for complex industrial fluids such as humid air.

Each approach is selected based on the characteristics of the fluid system in question and the specific demands of the simulation task. While emphasis is given to IDAES due to the larger context of the Ahuora Digital Twin platform, the approaches can fit with other simulation environments with similar features.

5.1.1 Equations of State: Illustration of complexity

An illustrative example of the complexity involved can be found in the implementation of cubic equations of state (EoS) within IDAES (Peng Robinson (PR)). The polynomial form of PR is shown in Equation 1. IDAES first uses a specialised cubic root finder to solve for Z given values for α , β , A and B . Once these roots are found and stored, the rest of the cubic EoS is written as Pyomo expressions (such as the departure functions for finding residual enthalpy, fugacity and entropy). A smooth vapour liquid equilibrium external function must also be created to handle the non-smooth region. This was a substantial undertaking from the IDAES creators.

$$Z^3 - (1 - B)Z^2 + (A - 2B - 3B^2)Z - (AB - B^2 - B^3) = 0 \quad \text{Equation 1}$$

$$A = \frac{\alpha ap}{R^2 T^2} \quad B = \frac{\beta p}{RT}$$

This chapter addresses these challenges through the design, implementation, and validation of advanced property packages tailored to the fluids and operating conditions relevant to the dairy sector, but readily extensible to other low- to medium-temperature industries. A three-pronged modelling strategy is adopted, with each approach selected according to the complexity of the fluid, the availability of experimental or reference data, and the computational performance requirements:

Beyond the creation of these models, the chapter examines their integration into the digital twin introduced in Chapter 4 and detailed in Chapter 6. Particular emphasis is placed on ensuring consistency between property packages when fluids interact within the same flowsheet, aligning reference states for enthalpy and entropy, and validating model outputs against both experimental data and industrial measurements.

The outcome of this work is twofold. First, it provides a set of high-fidelity, solver-compatible property packages capable of accurately representing the diverse fluids encountered in dairy process electrification studies. Second, it establishes a generalisable and reproducible workflow for property model development, enabling other sectors to adopt and adapt these methods without requiring specialist thermodynamic programming expertise.

5.2 The thermodynamic properties of milk

5.2.1 Introduction

Despite decades of dairy research and extensive experimental measurement of milk's thermophysical properties, there is currently no fully defined, general-purpose property package capable of calculating all required thermodynamic properties of milk across the wide range of conditions encountered in industrial processing. Numerous studies report specific properties such as density, heat capacity, viscosity, or boiling point elevation but these measurements are often limited to narrow composition ranges, specific temperatures, or fixed concentrations. As a result, process engineers must rely on piecing together disparate datasets and partial correlations, rather than drawing from a single, and internally consistent thermodynamic formulation.

Conventional process simulators, such as Aspen HYSYS, represent the de facto standard for industrial process modelling. These paid simulators contain extensive databases of fluids and built-in thermodynamic models, but these databases are primarily oriented towards petrochemical and related industries, with only limited coverage of aqueous electrolyte systems such as brine. Consequently, when modelling a milk evaporator in such environments, a custom property package must be developed. This typically involves selecting a base model such as a conventional activity coefficient formulation and introducing pseudo-components to represent the major milk constituents (water, fats, proteins, lactose, and mineral salts), together with binary interaction parameters.

5.2.2 Method for activity coefficient modelling for conventional simulation platforms

To model milk in Aspen HYSYS and DW Sim, a user must define pseudo-components, along with user-specified thermodynamic and physical property data. This methodology followed the compositional and structural assumptions presented by Zhang et al. (2014) who developed a simplified representation of milk based on major chemical constituents. In this framework, the fat component of milk was represented using palmitic acid and oleic acid, which are among the most abundant fatty acids in bovine milk. Minor fatty acids such as butyric, myristic, and stearic acid were excluded to simplify the model, based on their lower overall contributions to total fat content. The protein fraction was approximated using casein, which constitutes approximately 80% of total milk protein. Whey proteins, representing the remaining 20%, were omitted from the model due to their relatively low proportion and the lack of consistent data. The mineral content of milk, which includes calcium, sodium, potassium, magnesium, phosphate, citrate, and chloride ions, was further simplified in this model by representing the mineral phase as a binary mixture of sodium chloride and potassium chloride in equal proportions, again consistent with the assumptions of Zhang et al. (2014).

The molecular properties of each pseudo-component such as molar mass, density, and boiling point were defined using a combination of literature sources, including Zhang et al. (2014) and thermophysical property data reported by Munir et al. (2016), Hu et al. (2009), and Choi (1986). Table 6 summarises the final set of component properties used in the Aspen HYSYS simulation. In addition to standard properties, specific heat capacity data were incorporated using empirical polynomial coefficients for fat, protein, lactose, and salts (Table 7). These coefficients allow the simulator to adjust heat capacity values dynamically with temperature, thereby improving the accuracy of enthalpy calculations across the temperature range encountered in dairy processing operations.

Table 6 Final milk component properties used for process simulation (Zhang et al., 2014).

Properties	Water	NaCl	Casein	Oleic acid	Lactose	Total solids
Weight percentage	87.0	0.8	3.4	4.0	4.8	13
Molecular weight (kg/kmol)	18.02	58.44	23,000	282.5	342.30	–
Density (kg/m ³)	998	2165	1250	893.4	1525	–
Boiling point (°C)	100	1413	1000	358.9	668.9	–

Table 7 Specific heat coefficient values for fat, protein, lactose and salt used for simulation of raw milk in Aspen HYSYS (Munir et al., 2016)

Specific heat coefficients	Fat	Protein	Lactose	Salt	Average
A	1.98	2.01	1.55	2.93	2.21
B	1.47E-03	1.21E-03	1.93E-03	–	1.54E-03
C	-4.80E-06	-1.31E-06	-5.94E-06	–	-4.02E-06

The Equation of State (EoS) selected was NRTL. The Non-Random Two-Liquid (NRTL) model is an activity coefficient-based thermodynamic model widely used to describe the non-ideal behaviour of liquid-phase mixtures. It is particularly well-suited to systems with strong intermolecular interactions and compositional asymmetry, such as those involving polar, associating, or hydrogen-bonding species. Unlike cubic equations of state, which are primarily designed for vapour-phase and hydrocarbon systems, the NRTL model focuses on predicting liquid-phase activity coefficients by considering local composition effects. It assumes that the molecules within the liquid phase are not randomly distributed and introduces binary interaction parameters to represent the energetic differences between like and unlike molecular pairs (Renon & Prausnitz, 1968).

A key challenge in implementing the NRTL model for simulating milk lies in the limited availability of reliable vapour–liquid equilibrium (VLE) data for the relevant component pairs. Accurate VLE data are essential for determining the binary interaction coefficients that underpin the model’s ability to predict non-ideal phase behaviour. For many conventional chemical systems, these coefficients can be readily sourced from established databases or regression of experimental data. Experimental VLE data for binary pairs involving lactose, casein, or representative fatty acids (e.g., oleic or palmitic acid) are extremely scarce in the literature. This lack of data presents a significant barrier to fitting the NRTL model.

In the absence of sufficient experimental vapour–liquid equilibrium (VLE) data, group contribution methods such as UNIFAC (Universal Quasi-Chemical Functional-group Activity Coefficients) offer an alternative approach for estimating binary interaction coefficients (Fredenslund et al., 1975). Unlike models that rely solely on empirical data fitting, UNIFAC predicts activity coefficients based on the functional group composition of each molecule. Each compound is decomposed into its constituent chemical groups, and group interaction parameters derived from extensive experimental datasets are used to estimate the excess Gibbs energy and corresponding activity coefficients for mixtures. This method is particularly valuable when working with components for which direct VLE data are unavailable or impractical to obtain, as is the case species found in milk. One should note that, the applicability of UNIFAC remains limited by the extent of its group parameter library and its inherent assumptions of group additivity and ideal mixing within groups. While it may provide useful approximations for smaller components such as fatty acids or simple carbohydrates (e.g., lactose), its accuracy declines when applied to complex, high-molecular-weight species such as casein or aggregated protein structures. Therefore, while UNIFAC can serve as a practical estimation tool for initial model development and feasibility studies, further refinement or experimental validation is recommended for high-fidelity simulation of dairy process fluids.

For the simple chain fats and salts the molecular structure and function group information was accessible and simple to enter. A cruder estimation for Casien was given as seen in Table 8.

Table 8 List of Binary Interaction Coefficients for simulation of milk

Binary Interaction Coefficients	H2O	Oleic Acid	Sorbitol	Fatty Oils	Casein
H2O	–	8260.536	-3640.94	10373.15	10373.15
Oleic Acid	39381.79	–	20362.28	-1680.02	10373.15
Sorbitol	2619.551	24858.75	–	10373.15	10373.15
Fatty Oils	3403.502	6092.216	11374.63	–	10373.15
Casein	-1765.5	-1765.5	-1765.5	11374.63	–

The vapour phase was modelled using the ideal gas assumption, which is appropriate given the low-pressure conditions and the composition of the vapour stream (consisting of 99.9% water vapour). At the operating pressures typical of milk evaporation (~28 kPa), water vapour exhibits near-ideal behaviour, making the ideal gas approximation both computationally efficient and sufficiently accurate for energy and phase equilibrium calculations.

5.2.3 Method adaptation for equation orientated frameworks

Modelling the properties of milk within the IDAES framework due to its equation-oriented nature provided some additional challenges, such as continuous differentiability and the requirement of smooth functions. Fortunately, both an NRTL suite of equations and a VLE model with a smooth phase transition for activity coefficient models has been developed into the framework.

5.2.3.1 Explanation of smooth functions for equation-oriented simulation

The method used by IDAES is based on fugacity equality. Fugacity serves as a corrected pressure that accounts for deviations from ideal behaviour, providing a more accurate representation of phase behaviour in non-ideal systems.

Mathematically, the VLE condition for a component i can be expressed as:

$$f_i^{vap} = f_i^{liq} \quad \text{Equation 2}$$

where f_i^{vap} and f_i^{liq} represent the fugacity of component i in the vapour and liquid phases, respectively. Fugacity itself is not a directly measurable quantity but is related to more accessible variables through fugacity coefficients ϕ , which represent the ratio between the fugacity and the actual pressure. The expressions for fugacity in each phase are typically written as:

$$f_i^{vap} = y_i \phi_i^{vap} P \quad f_i^{liq} = x_i \gamma_i P_i^{sat} \quad \text{Equation 3}$$

where:

- y_i and x_i are the mole fractions in the vapour and liquid phases.
- ϕ_i^{vap} is the fugacity coefficient in the vapour phase.
- γ_i is the activity coefficient in the liquid phase.
- P_i^{sat} is the saturation pressure of pure component.
- P is the total system pressure.

Since the vapour phase is assumed to be ideal the fugacity coefficient is set to be 1 and from the previous work on milk, the UNIFAC parameter estimations can be used to calculate the activity coefficients.

In standard simulation practice, flash calculations are typically performed at a specified stream temperature. This approach is sufficient when the thermodynamic state of the stream is known in advance (is it fully liquid, fully vapour or a two-phase system) such as in sequential steady state simulation. However, in optimisation (equation orientated modelling), stream conditions are not always known by calculation and may vary as the solver analyses the solution space (Lee et al., 2021). As a result, the specification may fall outside the two-phase region, rendering the flash calculations/equations trivial or degenerate (i.e. it is totally in the liquid phase so the solution to the VLE equilibrium constraint is a moot point).

To address this issue, the property packages in IDAES that support vapour–liquid equilibrium (VLE) implement custom flash calculations at a calculated “equilibrium temperature.”

$$T_1 = 0.5 \left[T + T_{bubble} + \sqrt{(T + T_{bubble})^2 + \epsilon_1^2} \right]$$

Equation 4

$$T_{eq} = 0.5 \left[T + T_{dew} + \sqrt{(T + T_{dew})^2 + \epsilon_2^2} \right]$$

where:

- T_1 is an intermediate variable.
- T is the temperature of the stream.
- T_{eq} is the equilibrium temperature of the stream.
- T_{bubble} is the bubble point temperature of the mixture.
- T_{dew} is the dew point temperature of the mixture.
- ϵ_1, ϵ_2 are smoothing parameters.

The function is a smooth approximation for the maximum (in the case of T_1) and minimum (in the case of T_{eq}) values. These functions combine to give the following result. When the stream temperature is lower than the bubble point temperature, the vapour–liquid equilibrium (VLE)

calculations are performed using the bubble point temperature. Conversely, if the stream temperature exceeds the dew point temperature, the calculations are conducted at the dew point. In all other cases when the stream temperature lies between the bubble and dew points the VLE calculations are executed at the actual stream temperature. Using the smooth functions enables the problem to be differentiable thus allowing the use of gradient-based solvers.

This is a relatively simple example to illustrate the complex nature of gradient orientated solution methods and problem formation. For further examples in creating a smooth and square flash formulations for equation oriented flowsheet optimisation, read the paper published by Burgard et al. (2018).

5.2.4 Results for the milk activity coefficient models

For the formulation of mass and energy balances, accurate knowledge of heat capacity and boiling point data is essential. Heat capacity directly underpins enthalpy calculations and thus governs the reliability of energy balances, while boiling point information is fundamental to VLE analysis, particularly in flash calculations. In aqueous mixtures, where water is the primary solvent, boiling point elevation (BPE) is used to quantify the temperature difference between pure water and the solution at the same pressure, capturing the impact of dissolved solids on phase behaviour.

A key challenge in modelling milk is that available property methods are generally provided as separate empirical correlations for individual properties. For example, heat capacity correlations may be available as a function of temperature and composition, while BPE or VLE behaviour is described using a separate correlation developed from a different dataset. These correlations can be accurate within their fitted ranges, but they do not form a unified thermodynamic property framework. Consequently, when they are combined in a process simulator, there is no guarantee that heat capacity, enthalpy, concentration change, and VLE behaviour remain internally consistent across varying milk compositions. This limitation is particularly important for evaporation modelling, where water removal changes both the liquid composition and the boiling behaviour of the stream.

It is also important to note that in food systems such as milk, experimental property data can be uncertain, often varying by as much as $\pm 10\%$. This arises from the inherent complexity of these mixtures, where differing proportions and interactions of proteins, fats, and minerals across samples introduce significant variability.

Figure 21 compares the model predictions with experimental values (Lewis, 1993) for whole milk (4 wt% fat) over a temperature range of 0–65 °C. The model reproduces the general temperature dependence of the experimental data, with predicted values remaining within 2–3% of the reported measurements across most of the domain. Slight underprediction is observed between 10–20 °C, but overall agreement is strong, indicating that the pseudo-component representation is suitable for capturing bulk thermal properties at typical processing temperatures.

The slight underprediction observed in the 10–20 °C range is most likely attributable to the latent heat associated with the solid–liquid phase transition of milk fat. As the fat fraction undergoes melting within this temperature region, additional energy is absorbed that is not fully captured by the simplified modelling approach. In the present methodology, this effect was not explicitly incorporated, as the fat concentration in whole milk is relatively low and the operational relevance of this temperature range is limited. During processing, milk only exists in this state when handled as whole or skim milk, prior to concentration. Consequently, while the omission introduces a minor discrepancy at low temperatures, it does not significantly impact the overall applicability of the model for process integration and electrification studies.

Compared with the whole milk Minim et al. correlation, the proposed model gives lower heat capacity values and more closely follows the experimental data across much of the temperature range. The Minim et al. model shows a gradual increase in heat capacity with temperature and tends to overpredict the experimental data, particularly at higher temperatures. By contrast, the proposed model better represents the flatter experimental trend observed above approximately 20 °C. Some of these discrepancies maybe due to natural variations and oversimplifications of the heat capacity model.

Figure 22 presents the results for concentrated milk containing approximately 10 wt% fat (which corresponds to a total solids percentage by mass of 39%). As expected, the specific heat capacity of concentrated milk is lower than that of whole milk, reflecting the higher solids content. The model follows the experimental trend closely, with deviations of less than 2% across the 35–70 °C range (the temperature range evaporator effects and concentrate heaters operate within).




Figure 21 Comparison of Specific Heat Capacity (J/g·K) versus Temperature (°C) for whole milk (Composition of Fat equating to 4% by mass) against the model generated by Minim et al.(2002). Experimental data sourced from Lewis (1993).

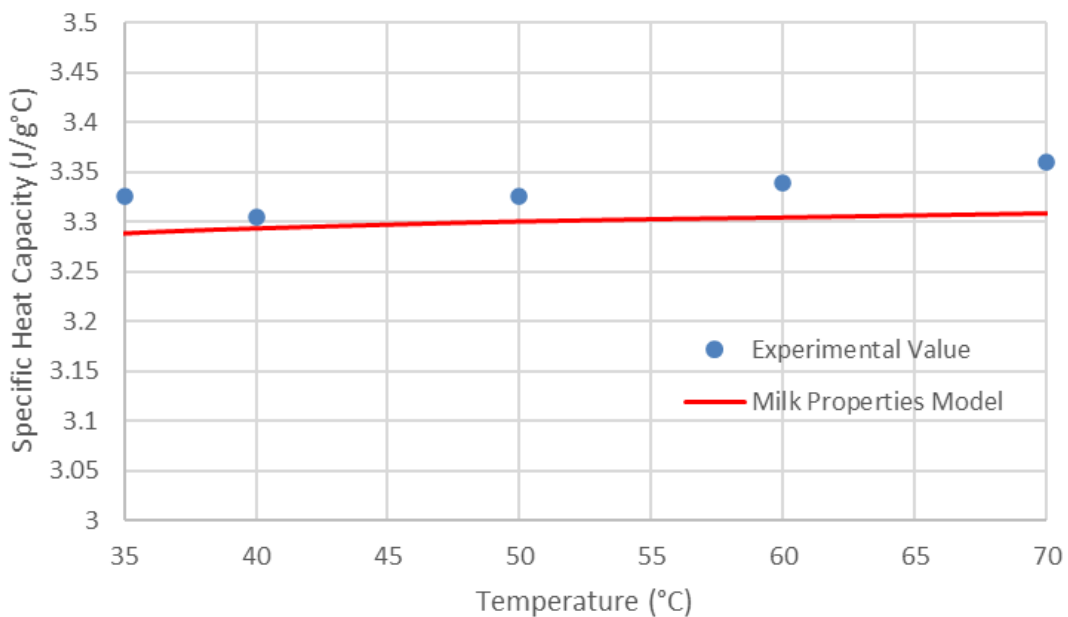


Figure 22 Comparison of Specific Heat Capacity (J/g·K) versus Temperature (°C) for concentrated milk (Composition of Fat equating to 10% by mass). Experimental data sourced from Lewis (1993).

The ability of the model to capture BPE with increasing dry mass is shown in Figure 23. The predicted curve aligns well with the experimental data and subsequent property model of Winchester (2000).

At higher solids contents (>40 wt%), the model slightly overestimates BPE, though the deviation remains within engineering tolerances and the data is sparse. The predicted trend confirms that the property package adequately represents the elevation of boiling temperature due to dissolved solids, a critical factor in evaporation modelling.

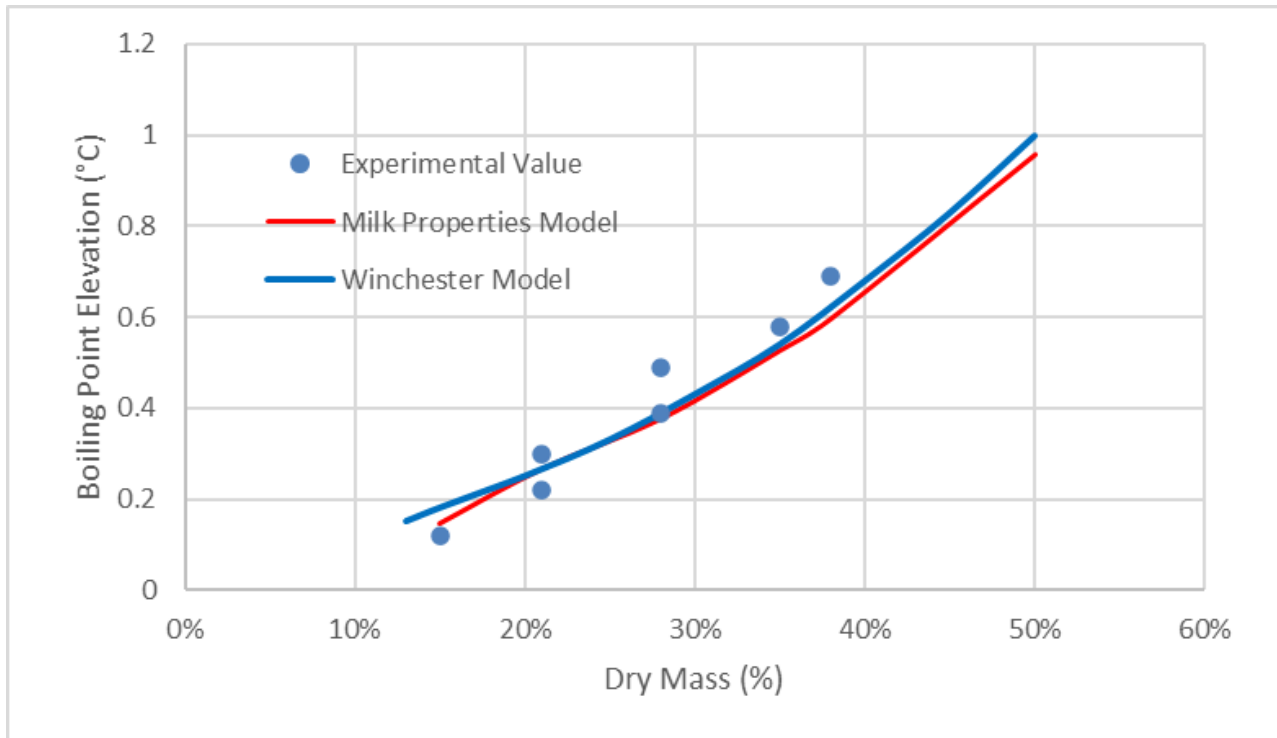


Figure 23 Comparison of Boiling Point Elevation (°C) versus Dry Mass (%) for a starting concentration of 90.7% water, 0.3% fat, 3.5% protein, 4.8% lactose, 0.7% minerals against the model generated by Winchester (2000). Experimental data sourced from Winchester (2000).

A subsequent test to assess the validity of the constructed pseudo-milk fluid, was using a constructed simplified process flowsheet consisting of a single heat exchanger. In the simulation scenario, 247 t/h of the pseudo-milk fluid was heated from 8 °C to 45 °C with an estimated a heat duty of 10,124 kW being calculated.

This result was benchmarked against the findings of Walmsley et al. (2016), who performed a similar heating operation using empirically measured milk properties derived from the work of Lewis (1993). Their reported heating requirement for the same process conditions was 10,049 kW. The resulting deviation between the two estimates was less than 0.75%, indicating good agreement and supporting the suitability of the pseudo-fluid model for thermal energy calculations. Furthermore, all simulated process stream temperatures for the base case were within 1 °C of the referenced design data, providing additional confidence in the model's applicability for energy-focused simulation tasks.

5.2.5 Discussion

For process engineers without specialised expertise in both mathematical programming and thermophysical property modelling, implementing accurate thermodynamic descriptions for complex or poorly documented fluids remains a significant barrier to applying the PI&E (Process Integration and Electrification) methodology. While the development of open-source property packages for such fluids offers a partial solution, the scalability is limited because it is impossible to create a property package that can be universally applied. Due to the diversity of process sites, each handling unique fluid compositions and formulations, this necessitates a simplified and accessible workflow that can be accessed by process and site engineers to rapidly and accurately develop a custom property package. One that enables practitioners to develop property models independently, provided that suitable data are available. A proposed method to support this objective is presented later in this chapter after discussing the implementation of highly accurate property packages for entropy calculation.

5.3 Custom Helmholtz-Based Formulations for Key Pure Components

5.3.1 Introduction

While standard property packages in commercial simulators are sufficient for conventional petrochemical systems, they often lack the accuracy and flexibility required for advanced applications, particularly when processes span wide ranges of pressure, temperature, and phase behaviour. To address this, more rigorous formulations are needed that can represent fluid properties to within experimental uncertainty and that are compatible with equation-oriented modelling environments.

Among the available approaches, Helmholtz-based equations of state have emerged as the most authoritative and widely adopted for high-accuracy property prediction. These formulations define the Helmholtz free energy as the fundamental potential, from which all thermodynamic properties (such as pressure, enthalpy, entropy, density, and speed of sound) can be rigorously derived using exact thermodynamic identities. This structure provides a mathematically consistent framework that enables simultaneous accuracy across vapour, liquid, two-phase, and supercritical regions, making it particularly well suited to the demands of modern energy systems, refrigeration cycles, and electrification technologies.

Having established the relevance and development of Helmholtz-based formulations in the literature, the focus of this section now shifts to their practical implementation within an open-source, equation-oriented simulation environment. The objective is not simply to reproduce existing capabilities found in tools such as REFPROP or CoolProp, but to create a property package that is fully compatible with the solver requirements of platforms like IDAES and the Ahuora framework. This requires a modelling approach that can deliver not only high-accuracy property predictions across multiple phases, but also the analytical derivatives necessary for robust convergence within large-scale optimisation problems.

The following subsections therefore describe the methods used to translate Helmholtz-based formulations into a form that can be embedded within a DT. This includes the generation of property expressions from the fundamental Helmholtz energy, the derivation and coding of the associated partial derivatives, and the validation of the resulting implementation against reference data. In this way, the work extends beyond using existing libraries and demonstrates how high-fidelity property models can be constructed and integrated directly into equation-oriented process simulation frameworks.

5.3.2 Method for creating a Helmholtz Free Energy Property Package for Equation Oriented Programming

The purpose of this section is to describe the method developed for implementing a Helmholtz-based property model directly within an equation-oriented simulation environment. As discussed in Chapter 4, equation-oriented platforms such as IDAES offer the flexibility required for advanced optimisation, process integration, and electrification studies, but they place additional demands on property models. In particular, they require not only accurate predictions of thermodynamic properties, but also the provision of first- and often second-order partial derivatives with respect to the chosen state variables to ensure numerical stability and solver convergence.

In contrast, most sequential-modular simulators do not natively expose Helmholtz-based formulations. Aspen HYSYS, for example, relies primarily on cubic equations of state and activity coefficient models. Although it does not embed Helmholtz equations, it can link externally to REFPROP or CoolProp via dynamic-link libraries (DLLs) to access high-accuracy thermophysical data. DWSIM, as an open-source sequential-modular simulator, already includes preconfigured CoolProp connections, allowing users to call selected Helmholtz-based models without additional

coding. These linkages extend the utility of sequential-modular tools, but they remain “black box” integrations where derivative information is not directly accessible (nor required) to the user or solver.

5.3.2.1 Helmholtz Free Energy

The method for evaluating dimensionless Helmholtz free energy is the summation of two components: an ideal-gas contribution and a residual (real-fluid) contribution, as shown in Equation 5. From this formulation, all relevant thermodynamic properties can be derived analytically using the first and second partial derivatives of the Helmholtz free energy with respect to the reduced density (denoted as δ) and the inverse reduced temperature (denoted as τ).

$$\frac{f(\delta, \tau)}{RT} = \phi^o(\delta, \tau) + \phi^r(\delta, \tau) \quad \text{Equation 5}$$

There are also a significant number of uncommon symbols utilised these are listed in Table 9

Table 9 Symbol for Equation 5

Symbol	Meaning	Unit
f	Specific Helmholtz free energy	kJ/kg
ϕ^o	Ideal dimensionless Helmholtz free energy	-
ϕ^r	Residual dimensionless Helmholtz free energy	-
δ	Reduced density, ρ/ρ^*	-
ρ^*	Reducing mass density, usually critical density	kg/m ³
τ	Inverse reduced temperature, T/T^*	
T^*	Reducing temperature, usually critical temperature	K

An example of such a calculation would be that of entropy (an important property for isentropic device calculations used in Heat Pumps). As mentioned previously to integrate a property into the equation-based framework at minimum the function value and its first derivative with respect to all

arguments must be given. Solving will be greatly improved if the second derivative with respect to all its arguments is also included.

The entropy function at non saturated conditions is a function of two variables δ and τ . The system of equations to calculate the functions value and its partial derivatives is shown below.

Intermediate equations:

$$z = \frac{\partial \phi^i}{\partial \tau} + \frac{\partial \phi^r}{\partial \tau} \quad z_\delta = \frac{\partial^2 \phi^i}{\partial \delta \partial \tau} + \frac{\partial^2 \phi^r}{\partial \delta \partial \tau} \quad \text{Equation 6, Equation 7}$$

$$z_{\delta\delta} = \frac{\partial^2 \phi^i}{\partial \delta^2} + \frac{\partial^2 \phi^r}{\partial \delta^2} \quad z_\tau = \frac{\partial^2 \phi^i}{\partial \tau^2} + \frac{\partial^2 \phi^r}{\partial \tau^2} \quad \text{Equation 8, Equation 9}$$

$$z_{\delta\tau} = \frac{\partial^2 \phi^i}{\partial \delta \partial \tau^2} + \frac{\partial^2 \phi^r}{\partial \delta \partial \tau^2} \quad z_{\tau\tau} = \frac{\partial^3 \phi^i}{\partial \tau^3} + \frac{\partial^3 \phi^r}{\partial \tau^3} \quad \text{Equation 10, Equation 11}$$

Binary functions to determine entropy:

$$f = R \cdot (\tau z - \phi^i - \phi^r) \quad \text{Equation 12}$$

$$f_{\delta\delta} = R \cdot \left(\tau z_\delta - \frac{\partial \phi^i}{\partial \delta} - \frac{\partial \phi^r}{\partial \delta} \right) \quad \text{Equation 13}$$

$$f_{\delta\delta^2} = R \cdot \left(\tau z_{\delta\delta} - \frac{\partial^2 \phi^i}{\partial \delta^2} - \frac{\partial^2 \phi^r}{\partial \delta^2} \right) \quad \text{Equation 14}$$

$$f_{\delta\tau} = R \cdot \left(z + \tau z_\tau - \frac{\partial \phi^i}{\partial \tau} - \frac{\partial \phi^r}{\partial \tau} \right) \quad \text{Equation 15}$$

$$f_{\delta\delta\tau} = R \cdot \left(z_\delta + \tau z_{\delta\tau} - \frac{\partial^2 \phi^i}{\partial \delta \partial \tau} - \frac{\partial^2 \phi^r}{\partial \delta \partial \tau} \right) \quad \text{Equation 16}$$

$$f_{\partial\tau^2} = R \cdot \left(2z_\tau + \tau z_{\tau\tau} - \frac{\partial^2 \phi^i}{\partial\tau^2} - \frac{\partial^2 \phi^r}{\partial\tau^2} \right) \quad \text{Equation 17}$$

In addition to entropy, thermodynamic functions for pressure, internal energy, enthalpy, Gibbs free energy, isobaric and isochoric heat capacities, specific volume, and speed of sound are also derived from the Helmholtz formulation. These property functions are implemented in the open-source IDAES framework and can be reviewed in the repository (https://github.com/IDAES/idaes-ext/blob/main/src/general_helmholtz/props.cpp). The IDAES platform includes functionality to wrap these functions as AMPL user-defined functions, allowing them to be accessed by AMPL-compatible solvers. This enables full integration with open-source modelling frameworks such as Pyomo and, by extension, the IDAES equation-oriented environment, facilitating the solution of thermodynamically consistent systems of equations.

5.3.2.2 Calculation of Helmholtz free energy derivatives

To access the complete range of thermodynamic property functions, it is necessary to implement both the ideal and residual components of the Helmholtz free energy, along with their first and second derivatives with respect to the reduced variables τ (inverse reduced temperature) and δ (reduced density). However, a key complication in this process is the lack of standardisation across the literature. Authors often adopt different symbols, functional forms, or basis conventions (mass-based vs. molar-based), making direct comparisons and reimplementing challenging. As a result, there is no universal Helmholtz formulation that can be generically applied across fluids without adaptation.

Ammonia is an example of this challenge. Gao et al. (2023) published a modern, high-accuracy Helmholtz formulation for ammonia, motivated by the difficulty conventional models have in capturing its strong intermolecular hydrogen bonding effects. This modern formulation (Equation 18 and Equation 19) is not currently included in the standard IDAES library and was therefore implemented and integrated into the platform by the researcher.

$$\alpha^0(\tau, \delta) = a_1 + a_2\tau + \ln \delta + (c_0 - 1) \ln \tau + \sum_{i=1}^3 v_i \ln[1 - \exp(-u_i\tau/T_c)] \quad \text{Equation 18}$$

$$\begin{aligned}
\alpha^r(\tau, \delta) &= \sum_{i=1}^5 n_i \delta^{d_i} \tau^{t_i} + \sum_{i=6}^8 n_i \delta^{d_i} \tau^{t_i} \exp(-\delta^{c_i}) \\
&+ \sum_{i=9}^{18} n_i \delta^{d_i} \tau^{t_i} \exp[-\eta_i(\delta - \epsilon_i)^2 - \beta_i(\tau - \gamma_i)^2] \\
&+ \sum_{i=19}^{20} n_i \delta^{d_i} \tau^{t_i} \exp\left[-\eta_i(\delta - \epsilon_i)^2 + \frac{1}{\beta_i(\tau - \gamma_i)^2 + b_i}\right]
\end{aligned} \tag{Equation 19}$$

The first challenge is implementing the first and second partial derivatives to be utilised by the solver. The reference paper did not publish these derivatives explicitly and therefore the researcher calculated said derivatives analytically with assistance from the differentiation tool Wolfram Alpha (Wolfram Research Inc, 2024).

Derivatives for the ideal portion of Helmholtz free energy:

$$\frac{\partial \phi^i}{\partial \delta} = \frac{1}{\delta} \tag{Equation 20}$$

$$\frac{\partial^2 \phi^i}{\partial \delta^2} = -\frac{1}{\delta^2} \tag{Equation 21}$$

$$\frac{\partial \phi^i}{\partial \tau} = n_0[2] + \frac{n_0[3] - 1}{\tau} + \sum_{i \in \text{rng}} \frac{n_0[i] \cdot g_0[i]}{T_c \left(\exp\left(\frac{\tau g_0[i]}{T_c}\right) - 1 \right)} \tag{Equation 22}$$

$$\frac{\partial^2 \phi^i}{\partial \tau^2} = \frac{1 - n_0[3]}{\tau^2} - \sum_{i \in \text{rng}} \frac{n_0[i] \cdot g_0[i]^2 \cdot \exp\left(\frac{\tau g_0[i]}{T_c}\right)}{T_c^2 \left(\exp\left(\frac{\tau g_0[i]}{T_c}\right) - 1 \right)^2} \tag{Equation 23}$$

$$\frac{\partial^2 \phi^i}{\partial \delta \partial \tau} = 0 \tag{Equation 24}$$

Derivatives for the residual portion of Helmholtz free energy:

$$\begin{aligned} \frac{\partial \phi^r}{\partial \delta} = & \sum_{i \in \text{rng}} n_i d_i \delta^{d_i-1} \tau^{t_i} + \sum_{i \in \mathcal{I}_2} n_i \delta^{d_i-1} \tau^{t_i} (d_i - c_i \delta^{c_i}) \exp(-\delta^{c_i}) + \\ & \sum_{i \in \text{rng}} n_i \delta^{d_i} \tau^{t_i} \left(\frac{d_i}{\delta} - 2a_i(\delta - e_i) \right) \exp(-a_i(\delta - e_i)^2 - b_i(\tau - g_i)^2) + \\ & \sum_{i \in \text{rng}} n_i \delta^{d_i} \tau^{t_i} \left(\frac{d_i}{\delta} - 2a_i(\delta - e_i) \right) \exp\left(-a_i(\delta - e_i)^2 + \frac{1}{b_i(\tau - g_i)^2 + b_i}\right) \end{aligned}$$

Equation
25

$$\begin{aligned} \frac{\partial^2 \phi^r}{\partial \delta^2} = & \sum_{i \in \text{rng}} n_i d_i (d_i - 1) \delta^{d_i-2} \tau^{t_i} + \sum_{i \in \text{rng}} n_i \delta^{d_i-2} \tau^{t_i} [(d_i - c_i \delta^{c_i})(d_i - 1 - c_i \delta^{c_i}) - c_i^2 \delta^{c_i}] \exp(-\delta^{c_i}) + \\ & \sum_{i \in \text{rng}} n_i \tau^{t_i} \exp(-a_i(\delta - e_i)^2 - b_i(\tau - g_i)^2) [-2a_i \delta^{d_i} + 4a_i^2 \delta^{d_i} (\delta - e_i)^2 - 4d_i a_i \delta^{d_i-1} (\delta - e_i) + d_i (d_i - 1) \delta^{d_i-2}] + \\ & \sum_{i \in \text{rng}} n_i \delta^{d_i-2} \tau^{t_i} [d_i^2 d_i (-4\delta a_i (\delta - e_i) - 1) 2\delta^2 a_i (2a_i (\delta - e_i)^2 - 1)] \exp\left(-a_i(\delta - e_i)^2 \frac{1}{b_i(\tau - g_i)^2 + b_i}\right) \end{aligned}$$

Equation
26

$$\frac{\partial \phi^r}{\partial \tau} = \sum_{i \in \text{rng}} n_i t_i \delta^{d_i} \tau^{t_i-1} + \sum_{i \in \text{rng}} n_i t_i \delta^{d_i} \tau^{t_i-1} \exp(-\delta^{c_i}) +$$

$$\sum_{i \in \text{rng}} n_i \delta^{d_i} \tau^{t_i} \left(\frac{t_i}{\tau} - 2b_i(\tau - g_i) \right) \exp(-a_i(\delta - e_i)^2 - b_i(\tau - g_i)^2) +$$

Equation
27

$$\sum_{i \in \text{rng}} \frac{n_i \delta^{d_i} \tau^{t_i-1} \left(b_i^2 t_i + 2b_i b_i t_i (\tau - g_i)^2 - (b_i t_i (\tau - g_i)^3 + 2\tau) \right) \exp\left(-a_i(\delta - e_i)^2 + \frac{1}{b_i(\tau - g_i)^2 + b_i}\right)}{(b_i(\tau - g_i)^2 + b_i)^2}$$

$$\frac{\partial^2 \phi^r}{\partial \tau^2} = \sum_{i \in \text{rng}} n_i t_i (t_i - 1) \delta^{d_i} \tau^{t_i-2} + \sum_{i \in \text{rng}} n_i t_i (t_i - 1) \delta^{d_i} \tau^{t_i-2} \exp(-\delta^{c_i}) +$$

$$\sum_{i \in \text{rng}} n_i \delta^{d_i} \tau^{t_i} \left[\left(\frac{t_i}{\tau} - 2b_i(\tau - g_i) \right)^2 - \frac{t_i}{\tau^2} - 2b_i \right] \exp(-a_i(\delta - e_i)^2 - b_i(\tau - g_i)^2) +$$

Equation
28

$$\sum_{i \in \text{rng}} n_i \delta^{d_i} \tau^{t_i} \left[\frac{4b_i^2(\tau - g_i)^2(2b_i(\tau - g_i)^2 + 2b_i + 1)}{(b_i(\tau - g_i)^2 + b_i)^4} - \frac{2b_i(-2g_i t_i + 2t_i \tau + \tau)}{\tau(b_i(\tau - g_i)^2 + b_i)^2} + \frac{t_i(t_i - 1)}{\tau^2} \right] \exp\left(-a_i(\delta - e_i)^2 + \frac{1}{b_i(\tau - g_i)^2 + b_i}\right)$$

$$\begin{aligned}
\frac{\partial^2 \phi^r}{\partial \delta \partial \tau} &= \sum_{i \in \mathcal{I}_1} n_i t_i d_i \delta^{d_i-1} \tau^{t_i-1} + \sum_{i \in \mathcal{I}_2} n_i t_i \delta^{d_i-1} \tau^{t_i-1} (d_i - c_i \delta^{c_i}) \exp(-\delta^{c_i}) + \\
&\sum_{i \in \mathcal{I}_3} n_i \delta^{d_i} \tau^{t_i} \left(\frac{d_i}{\delta} - 2a_i(\delta - e_i) \right) \left(\frac{t_i}{\tau} - 2b_i(\tau - g_i) \right) \exp(-a_i(\delta - e_i)^2 - b_i(\tau - g_i)^2) + \\
&\sum_{i \in \mathcal{I}_4} \frac{n_i \delta^{d_i-1} \tau^{t_i-1} (d_i - 2\delta a_i(\delta - e_i)) (b_i^2 t_i + 2b_i b_i t_i (\tau - g_i)^2 - b_i(\tau - g_i)(b_i t_i (\tau - g_i)^3 + 2\tau))}{(b_i(\tau - g_i)^2 + b_i)^2} \exp(-a_i(\delta - e_i)^2 \\
&\quad + \frac{1}{b_i(\tau - g_i)^2 + b_i})
\end{aligned}$$

Equation

29

The code implementation for writing these external functions is recorded in 11.1 Appendix One.

5.3.2.3 Calculation of the auxiliary saturated density curve

The auxiliary equations are used to estimate saturated liquid and vapour densities. Their primary role is to serve as robust initial guesses for the solution of the phase equilibrium (flash) equations. These auxiliary correlations offer sufficiently accurate estimates of saturation densities across a wide range of temperatures and pressures, thereby improving the convergence behaviour of nonlinear solvers. Since these equations only provide initialisation, calculation of their derivatives is not required.

$$\delta = c + \sum_{i=1}^h n_i \left(1 - \frac{T}{T_c}\right)^{t_i} \text{ for the liquid density (saturated)} \quad \text{Equation 30}$$

$$\delta = c \exp\left(\sum_{i=0}^h n_i \left(-\frac{T}{T_c}\right)^{t_i}\right) \text{ for the vapour density (saturated)} \quad \text{Equation 31}$$

5.3.2.4 Implementation of equation constants and coefficients

After encoding the external functions, the appropriate experimentally derived coefficients from Gao et al. (2023) had to be made available for the AMPL wrapper. This was done through a JSON file populated with the data shown in Table 10, Table 11, Table 12 and Table 13. The relevant code is shown in 11.1 Appendix One.

Table 10 Physical constants and characteristic properties of ammonia Gao et al. (2023)

Symbol	Property	Value
R	Molar gas constant	8.314462618 J mol ⁻¹ K ⁻¹
M	Molar mass	17.03052 g mol ⁻¹
T_c	Critical temperature	405.56 K
p_c	Critical pressure	11.3634 MPa
ρ_c	Critical density	13.696 mol dm ⁻³
T_{tp}	Triple-point temperature	195.49 K
p_{tp}	Triple-point pressure	6.05339 kPa
ρ_{tpv}	Vapor density at triple point	0.0037408 mol dm ⁻³
ρ_{tpl}	Liquid density at triple point	43.091 mol dm ⁻³

T_{nbp}	Normal-boiling-point temperature	239.834 K
ρ_{nbpv}	Vapor density at the normal-boiling point	0.052258 mol dm ⁻³
ρ_{nbpl}	Liquid density at the normal-boiling point	40.024 mol dm ⁻³
ω	Acentric factor	0.256
T_0	Reference temperature for ideal-gas properties	300 K
p_0	Reference pressure for ideal-gas properties	1.0 kPa
	Reference ideal-gas enthalpy at T_0	28989.81844833686 J mol ⁻¹
	Reference ideal-gas entropy at T_0 and p_0	159.53995277505977 J mol ⁻¹ K ⁻¹

Table 11 Coefficients of the ideal equation of state for ammonia Gao et al. (2023)

i	n₀	g₀
1	-6.59406093943886	
2	5.60101151987913	
3	4.0	
4	2.224	1646
5	3.148	3965
6	0.9579	7231

Table 12 Coefficients of the residual equation of state for ammonia Gao et al. (2023)

i	n_i	t_i	d_i	c_i	a_i	b_i	g_i	e_i	b_{ii}
1	0.006132232	1	4						
2	1.7395866	0.382	1						
3	-2.2261792	1	1						
4	-0.30127553	1	2						
5	0.08967023	0.677	3						
6	-0.076387037	2.915	3	2					
7	-0.84063963	3.51	2	2					
8	-0.27026327	1.063	3	1					

9	6.212578	0.655	1	0.42776	1.708	1.036	-0.0726	
10	-5.7844357	1.3	1	0.6424	1.4865	1.2777	-0.1274	
11	2.4817542	3.1	1	0.8175	2.0915	1.083	0.7527	
12	-2.3739168	1.4395	2	0.7995	2.43	1.2906	0.57	
13	0.01493697	1.623	2	0.91	0.488	0.928	2.2	
14	-3.7749264	0.643	1	0.3574	1.1	0.934	-0.243	
15	0.0006254348	1.13	3	1.21	0.85	0.919	2.96	
16	-0.000017359	4.5	3	4.14	1.14	1.852	3.02	
17	-0.13462033	1	1	22.56	945.64	1.05897	0.9574	
18	0.07749072839	4	1	22.68	993.85	1.05277	0.9576	
19	-1.6909858	4.3315	1	2.8452	0.3696	1.108	0.4478	1.244
20	0.93739074	4.015	1	2.8342	0.2962	1.313	0.44689	0.6826

Table 13 Coefficients for saturated density equations (Gao et al., 2023)

i	n_i, liquid	n_i, liquid
1	0.051236	-0.089966
2	3.7925	-3.8722
3	-3.5929	-8.1183
4	4.6409	-25.293
5	-1.9893	-54.279
6	1.5978	-400.83

5.3.2.5 Implementation of thermal conductivity and viscosity

For completeness both thermal conductivity and viscosity calculation were also added as auxiliary functions from the works Monogenidou et al. (2018) and Fenghour et al. (1995) respectively. While these calculations are important in the field of process engineering and process integration, they were not integral to the PI&E methodology presented and therefore can be reviewed in 11.1 Appendix One.

5.3.2.6 Implementation of other important refrigerants

A similar method was followed to integrate both n-butane and iso-butane, which again are widely regarded natural refrigerants for future decarbonisation work. The code and equations can be found in 11.1 Appendix One.

5.3.3 Results

Validation of the developed property packages for ammonia and butane is significantly more straightforward than that required for complex multicomponent fluids such as milk. As pure substances, both ammonia and butane have been extensively studied, with their thermophysical properties well documented across a broad range of temperatures and pressures (Gao et al., 2023), (Bücker & Wagner, 2006).

A key focus of the validation process was the accurate prediction of saturation properties, which are particularly challenging to compute due to the inherent discontinuities in thermodynamic gradients at the phase boundary. This makes them an ideal test of the robustness and fidelity of the implemented equation of state. To evaluate the model performance, multiple representations of the saturation curve (including pressure–enthalpy and temperature–entropy relationships) were generated using the property package. These were then compared against established reference data and literature correlations.




Figure 24 Pressure-Enthalpy relationship data derived from for Ammonia Gao et al. (2023)

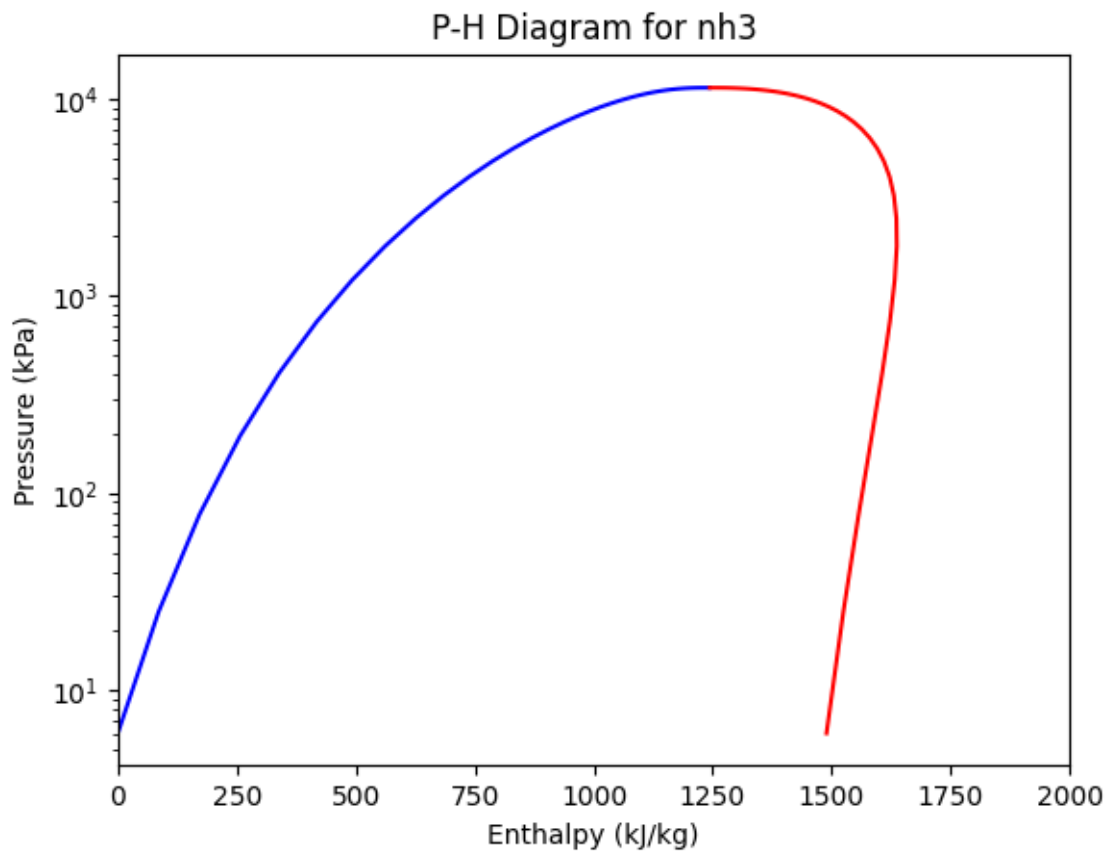


Figure 25 Pressure-Enthalpy relationship produced from generated package for Ammonia

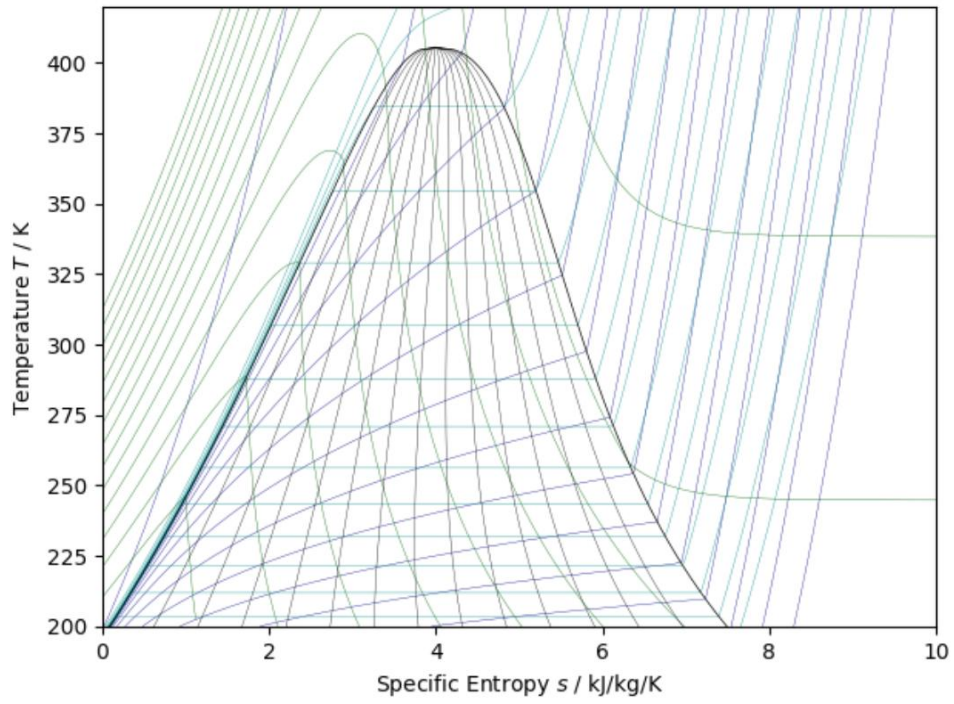


Figure 26 Temperature-Entropy relationship data derived from Gao et al. (2023) for Ammonia

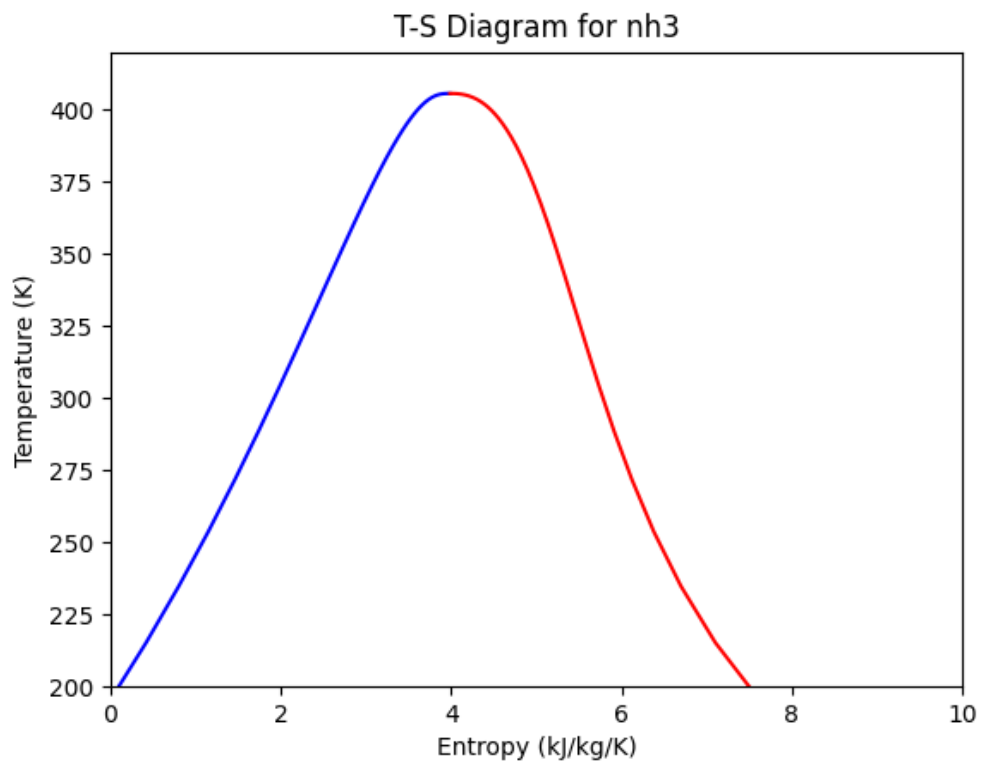


Figure 27 Temperature-Entropy relationship produced from the generated package for Ammonia

The comparison confirmed that the implemented models accurately reproduce the saturation envelope, capturing the expected behaviour near the critical point and exhibiting correct asymptotic trends in both the subcooled liquid and superheated vapour regions.

Another robust method for validating a thermophysical property package is to perform a point-by-point comparison between model-generated values and a verified reference property dataset. The comparison was done by directly benchmarking the output of the implemented equations against experimentally validated and/or internationally recognised sources. In this study, the National Institute of Standards and Technology (NIST) property database, specifically the NIST REFPROP system, was used as the primary reference for pure component fluids such as ammonia and butane.

The validation procedure involved generating a set of thermodynamic state points over a broad and representative range (the triple point (195.42 K, 6090 Pa) to 1.1 times the critical point (405.56 K, 11.357 MPa) of temperatures and pressures (the most common property pair a user would enter) at each state point, key properties such as specific enthalpy, entropy, and density were computed using both the implemented property model and the corresponding NIST reference data. Numerical comparisons were then performed, typically by calculating the absolute and relative percentage deviations for each property.




Figure 28 Error in Enthalpy Calculation vs Temperature and Pressure for Ammonia




Figure 29 Error in Entropy Calculation vs Temperature and Pressure for Ammonia




Figure 30 Error in Density Calculation vs Temperature and Pressure for Ammonia

An examination of Figure 28, Figure 29, and Figure 30 and its corresponding data, indicates that the developed property package demonstrates excellent agreement with the NIST reference model across the evaluated thermodynamic domain. In both the subcooled liquid and superheated vapour regions, the model achieved complete alignment with the reference data, exhibiting negligible deviation.

However, as the state points approached the critical region, minor discrepancies began to appear. The maximum observed deviation was approximately 0.4%, occurring at a temperature of 450 K and a pressure of 12 MPa. It is important to note that these states lie close to the critical point.

This inaccuracy near the critical point and above is to be expected. The Helmholtz EoS implementation used for Ammonia and the butanes followed the equations derived from the corresponding papers about said fluids (Bücker & Wagner, 2006; Gao et al., 2023). However whilst these equations are accurate a combination of numerical smoothing techniques, improper critical scaling and the highly sensitive nature of the critical region leads to discrepancies as shown by Bell et al. (2023). New literature has been published describing new tools used by NIST to vastly reduce solving time and ensure greater accuracy in the vicinity of the critical point through the use of Chebyshev expansions (Bell, 2024). One of the key methodological differences between the conventional EOS developers and the new method is the critical point is obtained by enforcing the constraint $\left(\frac{\partial p}{\partial \rho}\right)_T = \left(\frac{\partial^2 p}{\partial \rho^2}\right)_T = 0$, rather than imputing the experimental value. Other techniques such as the non-analytical Helmholtz parameter have also been used for critical point correction for selecting fluids like CO₂ and water (Span & Wagner, 1996).

Integration of high performance super ancillary curves into the Ahuora platform is interesting future work both for numerical smoothing and super critical accuracy. However, for the purposes of the PI&E methodology the accuracies achieved meet expectations due to the fact that very few systems that are not water and CO₂ go trans-critical.

5.3.4 Discussion

The successful implementation of custom Helmholtz-based formulations for ammonia and butane demonstrates the feasibility and value of integrating high-fidelity thermodynamic models into open-source, equation-oriented platforms such as IDAES. Again, the integration of these models into an open-source, equation-oriented simulation platform such as IDAES overcomes many of the long-standing limitations associated with proprietary software environments, including lack of transparency, extensibility, and solver compatibility for optimisation. These models enable precise calculation of properties across all fluid regions, supporting rigorous analysis for utility streams such as steam and electrical utility systems such as vapour compression cycles for the PI&E method. Validation against saturation curves and NIST reference data confirms their accuracy, with deviations remaining within engineering tolerances even near the critical point.

While this work establishes a robust framework for modelling pure fluids, broader adoption remains limited by complexity and the need for thermodynamic expertise. To address this, the following section proposes a simplified and generalisable workflow, aimed at enabling process engineers to independently develop and deploy property packages using accessible data and tools.

5.4 Machine-Learned Surrogate Models for Complex Mixtures

5.4.1 Introduction

The preceding sections have highlighted the technical complexity and specialised expertise required to implement accurate thermophysical property models within a fully equation-oriented framework, particularly for non-ideal fluids. While the development of preconfigured property databases for fluids such as milk and ammonia represents a meaningful advancement in improving access to Process Integration and Electrification (PI&E) methodologies, this approach alone is insufficient. Process engineers operate across a diverse range of industrial sites, each involving unique mixtures and fluid systems. As a result, it is impractical to provide a comprehensive library of pre-built models that adequately captures the full spectrum of fluid behaviour encountered in practice.

To address this limitation, an alternative approach is required. The method needs to retain the fidelity necessary for reliable process analysis, while significantly lowering the technical barrier to model development. Machine-learned surrogate models offer a promising solution to this challenge. By training on high-quality experimental data or outputs from high-fidelity simulations, these models can approximate complex thermodynamic behaviours without requiring explicit analytical expressions. Surrogates are capable of capturing non-linear, multivariate relationships inherent to non-ideal fluid systems, and once trained, can be deployed as callable, differentiable functions within equation-oriented environments.

In the context of PI&E, surrogate models enable rapid development of property packages tailored to specific site conditions or fluid compositions. They offer a scalable pathway for incorporating previously unmodelled fluids into integrated design and optimisation studies, thereby extending the reach and practical applicability of digitalised PI&E methodologies. The following section outlines the methodology adopted to construct, validate, and implement machine-learned surrogate models that can integrate into equation orientated frameworks, with a focus on maintaining consistency, lowering technical challenge, and solver compatibility.

5.4.2 *The case study: Humid Air*

To demonstrate the potential of machine-learned surrogate models within a PI&E context, this thesis presents the development and integration of a surrogate property model for humid air. Humid air is a ubiquitous working fluid in a wide range of thermal processes, and accurate thermodynamic representation is essential for modelling systems involving drying, condensation, air handling, and ventilation. Despite its prevalence, humid air exhibits non-ideal, multi-component behaviour, particularly at elevated temperatures, varying moisture contents and phase transitions, which complicates its accurate characterisation using traditional equations of state (Herrmann et al., 2009).

A particularly relevant application of humid air modelling arises in spray drying systems, a unit operation of critical importance in milk powder production. Spray dryers operate by contacting finely atomised liquid product with heated air in a controlled environment, rapidly removing moisture and transforming the liquid into powder. In such systems, the performance and efficiency of the dryer are strongly governed by the thermodynamic properties of the humid air stream specifically, its enthalpy and specific humidity.

Although the Ahuora platform provides flexibility for implementing custom property packages, the full humid air formulation presents a practical challenge for simultaneous equation-oriented simulation. The rigorous formulation involves large nonlinear expressions, phase-boundary logic, and iterative calculations for saturation and enhancement-factor behaviour. If these equations were embedded directly into the process flowsheet, they would substantially increase the number and complexity of equations passed to the solver, making the model more difficult to initialise and converge. This is particularly problematic when humid air properties are required repeatedly across multiple unit operations or operating states.

The humid air property model developed in this study was therefore motivated by the need to accurately represent the air–moisture mixture across the wide range of temperatures (typically 70–250 °C) and moisture loadings encountered in dairy spray dryers. Furthermore, because the spray dryer forms part of the end-to-end milk evaporation and drying process that serves as the overarching case study of this thesis, the model plays an integral role in extending digital PI&E analysis beyond evaporation alone, enabling a whole-process assessment of heat integration and electrification opportunities. For completeness the generated property model will also consider properties at and beyond the freezing point to add additional complexity and unlocking the ability to model refrigeration.

5.4.3 Method For Surrogate Modelling of Thermophysical Properties

Surrogate models are simplified mathematical approximations of complex functions or systems, constructed from input–output data. In the context of thermodynamic property modelling, a surrogate model replaces a computationally expensive or numerically challenging function (such as a detailed equation of state) with a data-driven regression model that retains accuracy within a defined operating domain. These models are particularly beneficial in optimisation contexts, where repeated function evaluations are required, and where symbolic derivatives are either difficult to obtain or prohibitively slow to compute.

To support the generation of such surrogates, this research utilised the Python Surrogate Modelling Objects (PySMO) library a dedicated open-source toolkit developed by the IDAES (Institute for the Design of Advanced Energy Systems) project. PySMO enables the automatic construction of surrogate models from simulation or experimental data using a suite of regression techniques shown in Forrester et al. (2008) (to which all mathematical equations are explained). These include:

- Polynomial regression models (linear, quadratic, or higher order),
- Radial basis function (RBF) models, and
- Kriging (Gaussian process) regression

Polynomial regression is the simplest and most interpretable surrogate modelling approach. It approximates the underlying data using a finite-order algebraic polynomial (for example a linear, quadratic, or cubic). These models are those simple for numerical solvers to interpret. However, polynomial regression lacks the flexibility to capture complex nonlinear behaviour in high-dimensional systems, particularly when the underlying physical phenomena are not easily approximated by low-order polynomials (Forrester et al., 2008).

Radial Basis Function (RBF) models address this limitation by constructing the surrogate as a weighted sum of basis functions (most commonly Gaussians). This approach allows for the accurate representation of highly nonlinear and multivariate functions, while maintaining smoothness and continuity in the model outputs and their derivatives. However, RBF models require careful selection of parameters such as the basis function width and can suffer from overfitting or poor extrapolation performance if not tuned appropriately (Forrester et al., 2008).

Kriging, or Gaussian Process Regression, is a probabilistic modelling technique that treats the target function as a stochastic process. It leverages a covariance structure defined over the input space to interpolate between observed data points with high accuracy. One of the key advantages of Kriging is its ability to quantify uncertainty in the model predictions, which can be especially valuable when data is scarce or expensive to generate. However, Kriging becomes computationally expensive as dataset size grows, due to the need to invert large covariance matrices, which limits its utility for extensive datasets (Forrester et al., 2008).

5.4.3.1 Creating a generalisable workflow for Hybrid Surrogate Models

As with all machine learning-based modelling approaches, the development of surrogate models involves a range of technical decisions, including the selection of hyperparameters, basis functions, data scaling strategies, and the division of training and validation datasets. These decisions can be highly specialised and present a significant barrier to process engineers aiming to adopt the PI&E methodology in practice. In the development of the surrogate model for humid air presented in this work, particular emphasis was placed on creating a generalisable and user-friendly workflow. The objective was to minimise the number of manual configuration steps required, thereby reducing the technical burden on end-users and facilitating broader adoption by practitioners without expertise in machine learning or advanced statistical modelling. These workflow decisions are described through this section.

5.4.3.2 Obtaining the data

Obtaining clean high quality data can be a challenge, conveniently Coolprop (Bell et al., 2014) has implemented a validated and rigorous humid air model that was proposed by Herrmann et al. (2009) which is callable via python scripting. A script was generated to create an expansive data set (over 20000 data points) across the whole validated range (130 K to 623.15 K, 0.01 kPa to 10 MPa, and humidity ratios up to 10 kg_w/kg_a). The entire script is transcribed in Appendix Two: Code snippets for surrogate property package modelling for readers to generate the data themselves. Reader should also note that 20000 data points is not required to generate the surrogate model (it is usually non optimal for performance to use such a large quantity).

5.4.3.3 Determining modelling variables

Conventionally when machine learning users create models from data, the user will determine the inputs required to run the model and the outputs generated from the model. That method is also true

for the creation of the surrogate model. However, as the number of input or output functions increases in a machine learning model, two critical issues commonly arise. The first is data sparsity (Abo-Sinna et al., 2014). As dimensionality (the number of variables) grows, the volume of the input space increases exponentially. As a result, an exponentially increasing amount of training data is required to adequately represent the full input space and capture meaningful patterns. Without sufficient data coverage, the model struggles to generalise effectively, leading to poor predictive performance.

Secondly, overfitting and computational cost increase substantially. Higher-dimensional data tends to exacerbate overfitting, as models can inadvertently fit noise rather than signal as they require more data to learn the meaningful patterns. The need for more data places greater computational burdens on training algorithms, therefore requiring greater and greater computational resources.

These challenges in high-dimensional machine learning models have been the subject of extensive research. Techniques such as Principal Component Analysis (PCA) and regularization have been developed to mitigate the dimensional challenges. While these techniques improve performance, an additional strategy can be employed to further reduce the complexity and training burden of the machine learning model

In the context of how the thermophysical property models are expressed in this research, both physics-based and data-driven models are ultimately composed of equations. This presents an opportunity for the user to selectively allocate the modelling of certain phenomena to either the physics-based or data-driven components, depending on appropriateness and complexity. For instance, it is inefficient and unnecessary for a machine learning model to "learn" fundamental physical laws such as mass conservation, unit conversion between moles and mass, or the summation rules that underpin partial enthalpy and entropy contributions. These are better handled explicitly using deterministic equations, allowing the surrogate model to focus on capturing more complex or less well-defined relationships within the thermodynamic system. These relationships already exist in the property models discussed previously.

An additional opportunity to reduce the dimensionality of surrogate models arises through careful partitioning of the thermodynamic system based on phase behaviour and modelling complexity. Using humid air as an illustrative example, the vapour phase is inherently complex, involving non-ideal interactions, entrainment phenomena, and temperature-dependent mixing behaviours. In contrast, once condensation has occurred, the condensed phase can be well approximated as pure liquid water under typical process conditions. Given the extensive research and high-fidelity

modelling of water properties (such as the internationally adopted IAPWS-95 formulation seen in the Helmholtz section) it makes sense to represent the liquid phase using established physics-based equations rather than the surrogate model learning this well-characterised behaviour.

This selective allocation of model responsibility not only reduces the surrogate model's dimensional and computational load but also enhances overall model reliability. Whenever a portion of the fluid's behaviour can be described by robust, validated thermophysical models, these should be prioritised within the physics-based framework. The approach can be generalised to other systems where phase separation occurs and where at least one component, for example water in many evaporation and drying processes.

By applying the described dimensional reduction techniques, it becomes possible to identify the minimum set of variables that must be modelled using data-driven methods in order to accurately represent a fluid system. In the case of humid air, the essential input variables are its thermodynamic state descriptors: temperature, pressure, and absolute molar composition which is typically expressed as the mass or mole fraction of water vapour (e.g., kg H₂O per kg total mixture). This selection is consistent with the Gibbs phase rule, expressed as $F = C - P + 2$, where F is the number of degrees of freedom, C is the number of components, and P is the number of phases. When treating dry air as a single pseudo-component, the number of degrees of freedom becomes three ($F = 2 - 1 + 2$), confirming that only three independent input variables are required.

For Process Integration and Electrification (PI&E) applications, the minimum required output variables are those that define the key thermodynamic properties of the vapour phase or the vapour liquid equilibrium. These include: the maximum water vapour content at saturation conditions for a given pressure and temperature (x_{wsat}), the molar enthalpy of the vapour phase (h_{gas}), the molar entropy of the vapour phase (s_{gas}), and the molar specific volume of the vapour phase (v_{gas}). These outputs provide the necessary thermodynamic description for integration studies involving humid air. All other relevant properties such as those related to the liquid or solid phase of water, or mass-based conversions can be calculated using existing, well-established physics-based equations.

5.4.3.4 Simple training method for a surrogate model

After deciding the variable required a surrogate model can be trained. This can pose a challenge to users with a low familiarity with the topic. One of the foremost decisions is the selection of the kernel function which can range from a Gaussian, multiquadric, or inverse multiquadric. Each of which defines how influence is distributed from each data point across the input space. Another key aspect

is the scaling and normalisation of the input features, as RBF models are sensitive to the relative magnitude of input variables. Additionally, the choice of training data points and their distribution across the domain is critical, as uniformly or adaptively sampling the input space can affect performance, especially in high-dimensional regimes (Forrester et al., 2008).

To assist a user multiple techniques were studied accessing their impact on accuracy and computational speed to determine a general workflow that could be used.

5.4.4 Results for simple surrogate model

The most straightforward method for improving the accuracy of surrogate models is to increase the number of training data points. To evaluate this effect, surrogate models based on the three common regression techniques (polynomial regression, radial basis function (RBF) networks, and kriging (Gaussian process regression)) were tested using datasets of 100, 1,000, and 10,000 randomly sampled input-output pairs. For the RBF models, a Gaussian kernel was employed with regularization to improve generalization and prevent overfitting. The polynomial models included terms up to fourth order, while the kriging models again utilised regularization as well as a Basinhopping solution method. The trained models are subsequently tested versus a new set of 20000 data points not seen in training.

Table 14 Surrogate model statistics across a range of models and data-points

Technique/Data point	Training Time/ Testing time	R ²	RMSE
Polynomial 100 data-points	24 s 10s	x_w_sat:0.86	x_w_sat:0.087
		h_gas:0.822	h_gas:1437
		s_gas:0.834	s_gas:4.645
		v_gas:0.959	v_gas:0.00192
Polynomial 1000 data-points	40 s 10s	x_w_sat:0.922	x_w_sat:0.0665
		h_gas:0.859	h_gas:1280.7
		s_gas:0.875	s_gas:4.036
		v_gas:0.987	v_gas:0.0011
Polynomial 10000 data-points	55 s 10s	x_w_sat:0.929	x_w_sat:0.063
		h_gas:0.865	h_gas:1251.49
		s_gas:0.881	s_gas:3.941
		v_gas:0.988	v_gas:0.001

RBF	4 s	x_w_sat:0.81	x_w_sat:0.102
100 data-points	15 s	h_gas:0.939	h_gas:841
		s_gas:0.938	s_gas:2.83
		v_gas:0.974	v_gas:0.0015
RBF	1 min 40 s	x_w_sat:0.929	x_w_sat:0.063
1000 data-points	9 mins	h_gas:0.865	h_gas:1251.49
		s_gas:0.881	s_gas:3.941
		v_gas:0.988	v_gas:0.001
RBF	5h 10 min 13 s	x_w_sat:0.999	x_w_sat:0.00061
10000 data-points	22 mins	h_gas:0.999	h_gas:101.7
		s_gas:0.999	s_gas:0.325
		v_gas:0.999	v_gas:1.49e-5
Kriging	50 s	x_w_sat:0.984	x_w_sat:0.0298
100 data-points	2s	h_gas:0.951	h_gas:752
		s_gas:0.935	s_gas:2.913
		v_gas:0.998	v_gas:0.003
Kriging	50 min 37 s	x_w_sat:0.999	x_w_sat:0.0012
1000 data-points	2s	h_gas:0.997	h_gas:177.24
		s_gas:0.997	s_gas:0.539
		v_gas:0.999	v_gas:4.08e-5
Kriging	Not Attempted	Not Attempted	Not Attempted
10000 data-points			

Table 14 shows that polynomial regression showed a steady increase in R^2 and a reduction in RMSE with larger datasets, achieving R^2 values above fair accuracy with 10,000 points. However, its performance for enthalpy and entropy remained lower (0.86 R^2), indicating limited representational capacity for these nonlinear outputs despite increased data.

RBF models demonstrated superior predictive capability for even at lower data volumes. At 100 data points, RBF outperformed polynomial regression for most variables except x_{sat} . With 10,000 data points, the RBF model achieved near-perfect agreement with ground truth data (0.999 R^2 , RMSE < 1% for all variables). This accuracy, however, came at a significant computational cost, with training time exceeding five hours and testing time extending to 22 minutes. This is typically due to the large and usually ill-conditioned matrix of equation that RBF generates (this is usually from oversampling) (Majdisova & Skala, 2018). There is also the risk of overfitting however due to completeness of the data it trained on this is less of a concern.

Kriging provided a favourable balance of accuracy and computational cost at smaller data sizes. With just 100 training points, the model achieved a minimum R^2 of 0.93 across all outputs, outperforming

both polynomial and RBF models at the same data level. At 1,000 points, Kriging reached parity with the high-resolution RBF model in a minimum accuracy of $0.997 R^2$, while requiring only 50 minutes of training and 2 seconds of evaluation. Due to the exponential scaling of computational complexity, Kriging was not attempted at 10,000 data points. Another metric worth mentioning is the models file size. Due to the fact Kriging uses weighted linear combination of all the training data points added, file sizes exponentially increase with data points (along with training time). For example, the 1000 data point file is 6.5 MB compared to polynomials 5 KB and RBFs 118 KB at the same number of training points. This can be expensive if a large set of surrogates were needing to be stored for a platform.

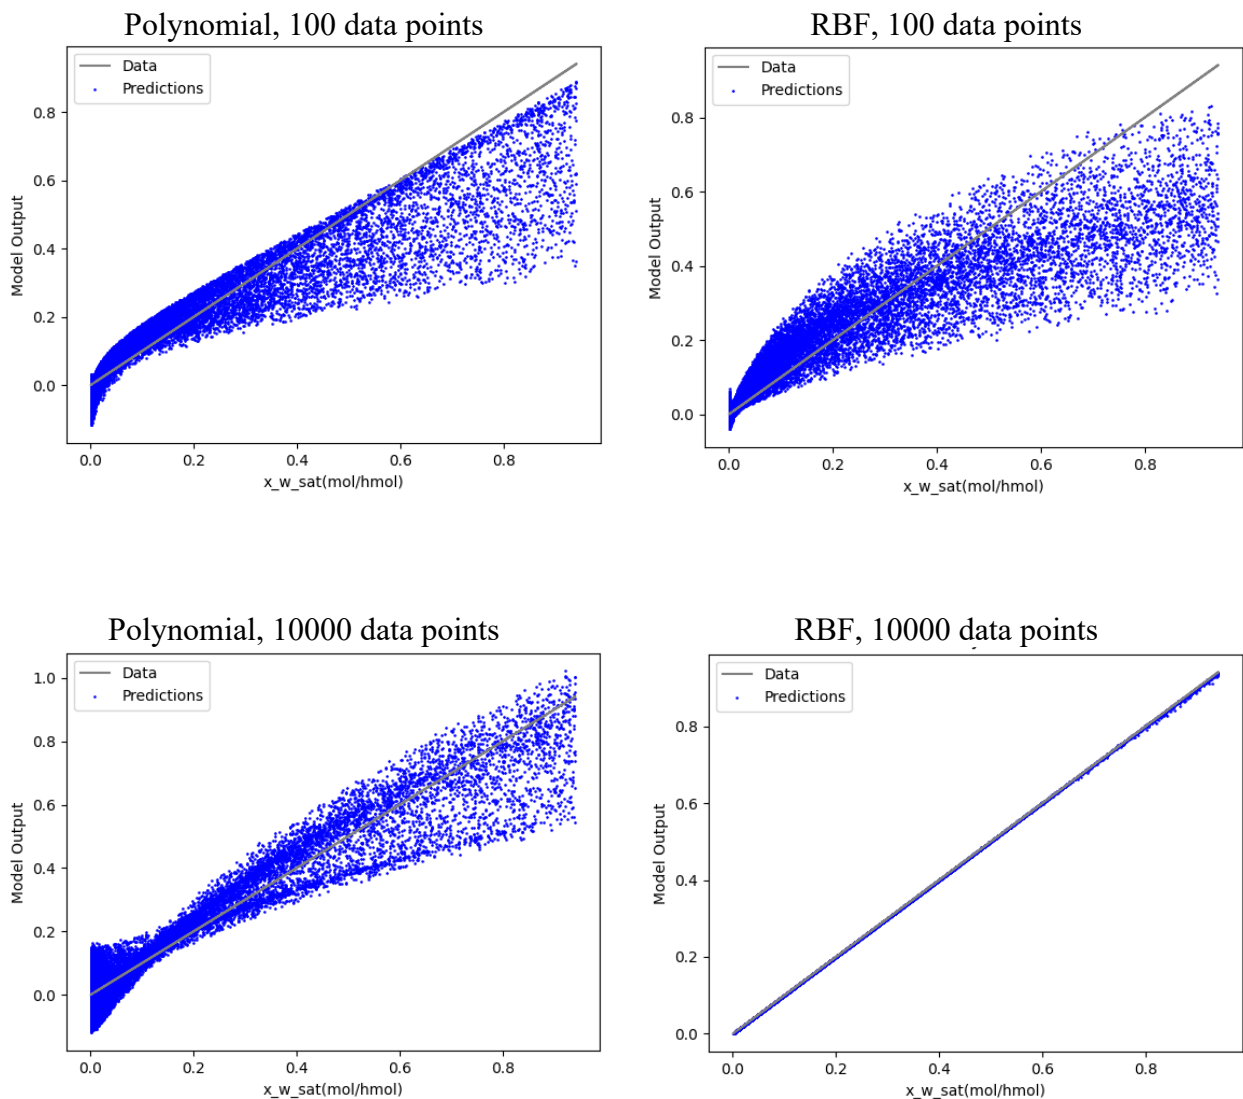


Figure 31 Parity plots showing how a polynomial basis function fails to model the nonlinear system even with a significant number of training points.

RBF, 1000 data points

Kriging, 1000 data points

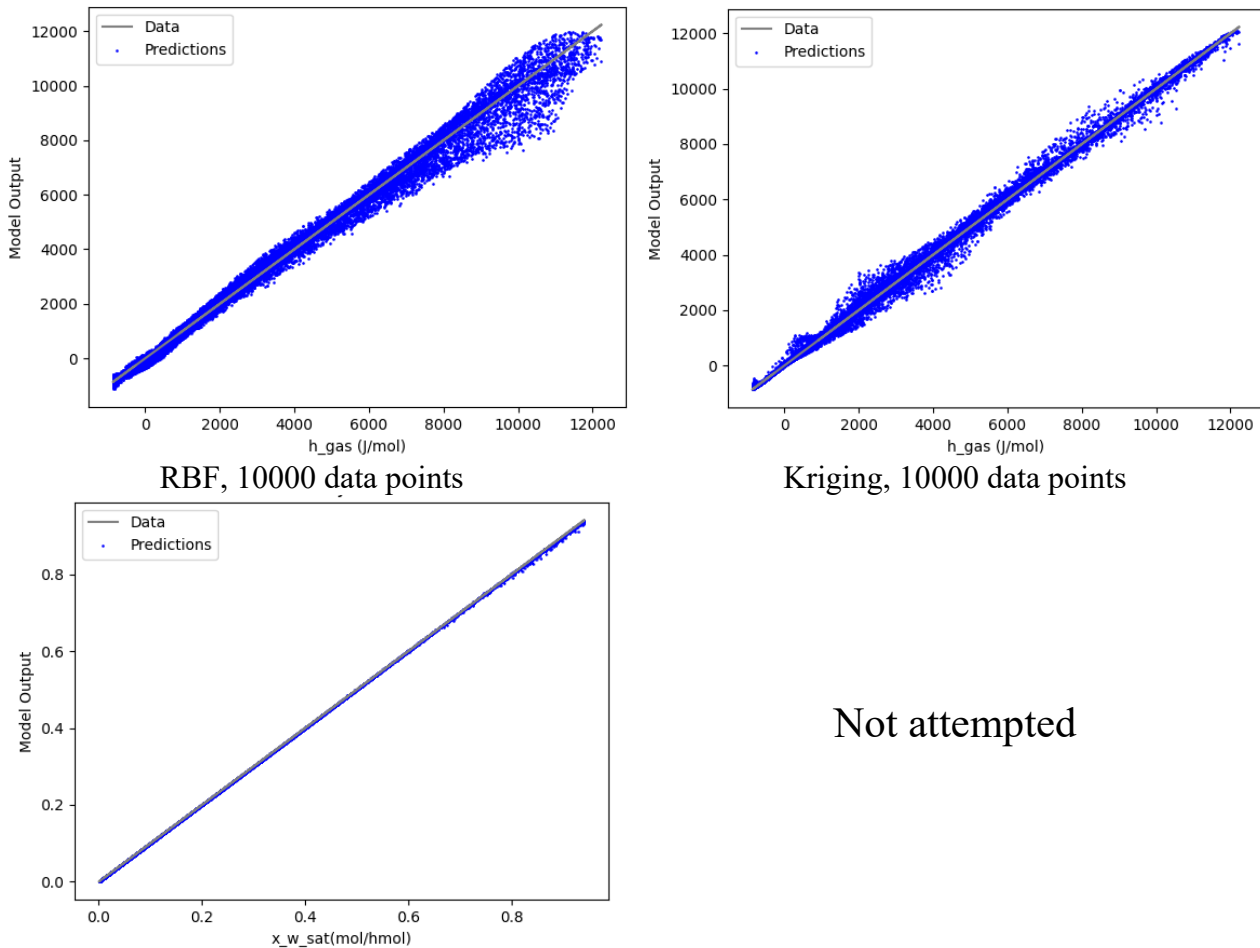


Figure 32 Parity plots for RBF and Kriging surrogates for h_{gas}

5.4.5 Method for increasing surrogate performance metrics via sampling

Whilst increasing the amount of training data does improve model accuracy across the board it also creates other trade-offs for users to consider. Currently the Kriging performed the most optimal out of the three being accurate and quick to evaluate. Its long training time is a heavy up-front cost to pay however it only need to be paid once.

However multiple literature studies have also identified the relationship between model complexity and training sample size (Forrester et al., 2008), (Mckay et al., 2000) suggesting that a subsampling routine can reduce complexity, whilst maintaining model accuracy. Multiple advanced “one-shot” techniques such as Latin Hypercube (Mckay et al., 2000) and Hammersley sampling (Wong et al., 1997) have been proposed in the past. The goal of these techniques is to obtain a more even distribution of sample points within the solution space, which in theory makes each data point more “valuable”.

Whilst these techniques are valuable for smooth functions the challenge in developing surrogate models for thermodynamic properties lies in addressing the non-smooth areas of the solution space. These take the form of both the non-uniform nonlinear behaviour and the discontinuities inherent in physical systems. Unlike smooth mathematical functions, thermodynamic property models often exhibit complex, region-specific nonlinearity, particularly in the vicinity of phase transitions or under extreme input conditions.

A second, more critical issue is the presence of physical discontinuities at phase boundaries. For example, in the case of humid air, the system displays a discontinuous shift at the phase change between liquid water and ice. Approaching the freezing point from the liquid side may yield a slightly different temperature than approaching it from the solid phase, leading to a small but distinct gap in predicted values (e.g., 0.01 °C vs. -0.01 °C). While such discontinuities are handled effectively in physics-based Helmholtz models, they present a major obstacle for surrogate models, which inherently attempt to construct a single smooth function. As a result, surrogate models often underperform near phase boundaries, over or underestimating values in an attempt to interpolate across the discontinuous region.

Figure 33 and Figure 34 show this behaviour clearly. Both are 3-D plots for all the input variables with bright/dark colouring showing significant deviation from the testing data. The large yellow region in Figure 33 is directly centred around the freezing point of water. The RBF model in its attempt to interpolate across the non-continuous region have substantially over predicted. Conversely Kriging due to its different approach to interpolation models this discontinuity better, however it deviates significantly in areas of low water concentration (particularly at atmospheric conditions). This shows the regional nonlinear effect real fluid exhibits.




Figure 33 RBF model for 100 data points of the humid air system




Figure 34 Kriging model for 100 data points of the humid air system

A uniform sampling approach could potentially miss these problem areas. Also, different fluids at different conditions will have different regions of non-smooth behaviour making manual or targeted sampling ineffective as well.

Recent literature surrounding water chemistry has proposed that adaptive sampling could potentially solve the issue surrogates face in computing non-smooth functions (Amusat et al., 2024). They proposed and showed that adaptive sampling would allow models to achieve higher accuracies than uniform distribution techniques of the same data size by strategically placing training samples where they are most needed to improve the model. Each adaptive sample is selected solely on its value added to the model's accuracy. The paper recommended the RBF framework for their water chemistry models however due to Kriging's performance previously the researcher will test both frameworks. Polynomial was discounted due to its inability to model the behaviour.

5.4.5.1 Adaptive sampling method

This sudden adapts the adaptive sampling method proposed by Amusat et al. (2024), converting it for its application to fluid properties. As mentioned, previous the surrogate models are build built using the PySMO framework and will use metric and sampling tools from said library.

1. Base surrogate model created from uniform sampling

The surrogate model construction begins with the generation of an initial training dataset using a Hammersley sequence. This method ensures a uniform and representative sampling of the input domain, avoiding the bias and clustering often associated with purely random sampling.

2. Classification accuracy improvement sampling

The classification-focused sampling phase serves as the first adaptive refinement step following initial surrogate model training. Specifically, once the initial surrogate is trained, it is evaluated across the broader dataset, and predictions are compared against actual data labels to determine false positives and false negatives. The algorithm then isolates these misclassified points and computes the absolute prediction error for each. The data point exhibiting the highest deviation between predicted and actual values is appended to the training set, and the model is retrained. This targeted approach ensures that new data points provide the most value for improving the model's decision boundary, especially in critical operational regions, and helps reduce classification error with minimal data augmentation.

3. Accuracy refinement sampling

In the second adaptive sampling phase, the focus shifts from classification to accuracy refinement within a specific and operationally important region of the solution space. Within this subdomain,

which is often critical for engineering decision-making, the model's predictions are evaluated across all unused data points, and the point exhibiting the maximum absolute error is identified. This data point is then added to the training set, and the model is retrained accordingly. This iterative process continues until a predefined number of additional samples have been selected.

4. The final model is returned and saved

For presentation clarity, only the results for the most non-linear and discontinuous variable, h_{gas} are shown in the proceeding tables and figures.

5.4.6 Results for adaptive sampling surrogate model creation

The adaptive sampling method was evaluated using the same datasets employed in the previous surrogate modelling trials. For each model, an initial training set of 50 uniformly distributed data points was generated using the Hammersley sequence sampling technique. Subsequent trials were conducted by augmenting this base dataset with either 50 or 100 adaptively selected data points. Because adaptive sampling involves retraining the surrogate model after each new point is added, an additional comparison was performed to assess how the size of the candidate sampling pool (from which new points are selected) impacts both model accuracy and computational performance. This analysis aimed to determine the trade-offs between pool size, sampling efficiency, and final model fidelity.

Table 15 Adaptive sampling model statistics across a range of models and data-points for h_{gas}

Technique/Data point	Training Time/ Testing time	R ²	RMSE
RBF (20000 pool) 50 adaptive data-points	28 min 19 s 9 s	$h_{gas}:0.993$	$h_{gas}:281$
RBF (5000 pool) 50 adaptive data-points	2 min 36 s 9 s	$h_{gas}:0.993$	$h_{gas}:274$
RBF (5000 pool) 100 adaptive data-points	7 min 18 s 14 s	$h_{gas}:0.996$	$h_{gas}:222$
Kriging (20000 pool) 50 adaptive data-points	56 s 2 s	$h_{gas}:0.995$	$h_{gas}:229$
Kriging (5000 pool) 50 adaptive data-points	38 s 2 s	$h_{gas}:0.995$	$h_{gas}:224.7$
Kriging (5000 pool) 100 adaptive data-points	6 min 45 s 2 s	$h_{gas}:0.997$	$h_{gas}:185.6$

The RBF model demonstrated excellent performance across all configurations, achieving an R^2 value of 0.993 with only 50 adaptively selected data points, regardless of whether the sampling pool size was 5,000 or 20,000. However, the larger 20,000-point pool required significantly more training time (28 minutes vs. 2.6 minutes) due to the increased computational burden of evaluating each candidate point during adaptive sampling. Despite this additional cost, the larger pool did not yield any appreciable improvement in model accuracy. This finding suggests that once a sufficiently diverse candidate pool is established, further expansion offers diminishing returns in terms of model performance.

Kriging also exhibited consistently strong predictive accuracy. With 50 adaptive samples, both the 5,000 and 20,000 pool sizes yielded an R^2 of 0.995, and RMSE values of 224.7 kJ/kmol and 229 kJ/kmol, respectively. When the adaptive sample count increased to 100, the model further improved to an R^2 of 0.997 and an RMSE of 185.6 kJ/kmol.

Both surrogate modelling techniques significantly benefited from adaptive sampling, with the technique substantially reducing the number of training point required for high accuracy. The reduction in required data subsequently improved the training time for whilst also improving the evaluation time for the RBF model. However, the results still underscore that Kriging offers a better balance between accuracy and computational efficiency, especially in smaller or moderately sized candidate pools. The RBF model, while slightly more accurate with increased data, incurs higher computational cost as sample counts and pool sizes grow. Therefore, for applications constrained by training time or computational resources, Kriging may be the preferred method, while RBF offers a marginally higher fidelity model when greater computational investment is acceptable.

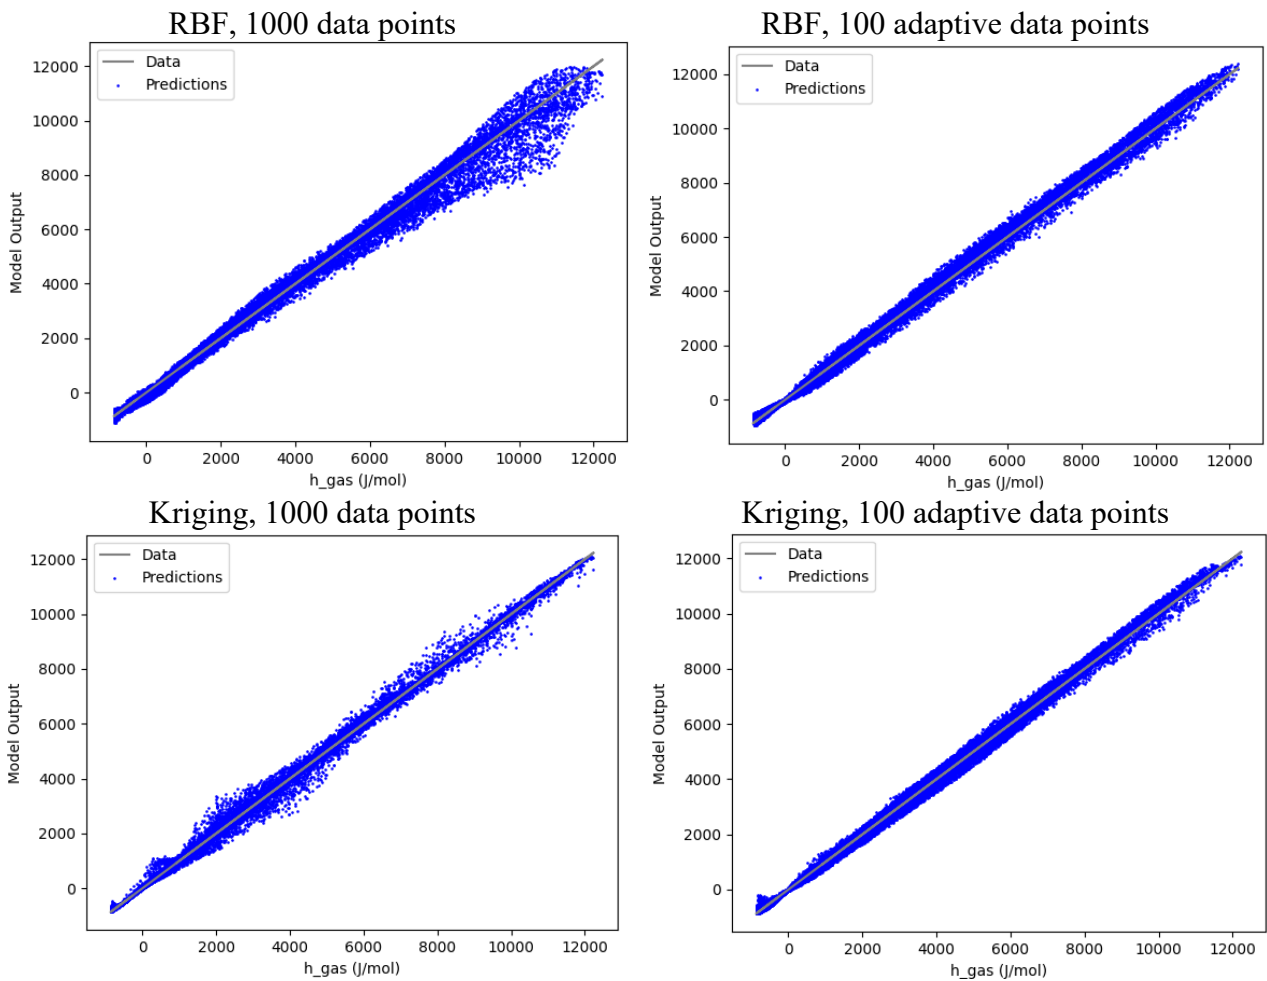


Figure 35 Parity plots to compare adaptive sampling for RBF and Kriging surrogates

Figure 37 shows that across the entire solution space Kriging is accurate, apart from the freezing point region which accounts for the majority of the error. This error is unavoidable even with a placing a significant number of adaptive training points at the freezing point. Again, this is due to the discontinuity present in this region. For the application of spray drying this error is acceptable and therefore the model is judged to be a success. The model will also be highly acceptable for deep refrigeration and HVAC as these applications rarely require exact solution at the freezing point.




Figure 36 Percentage difference in h_{gas} for the RBF surrogate model




Figure 37 Percentage difference in h_{gas} for the Kriging surrogate model

5.4.7 Creating the full property package

Upon creating the surrogate model, the complex phenomena, the additional variable and dimension removed at the beginning are then added back via simple equations. This implementation follows how the NRTL model for milk was implemented however the binary interaction coefficients are removed and the surrogate models is injected. Another addition was instead of writing the property rules for pure water in the liquid phase the highly accurate Helmholtz model was also injected (Figure 38).

```
1. def _make_state_vars(self):
59.
60.     inputs = [self.temperature, self.pressure, self.mole_frac_comp["water"]]
61.     outputs = [self.mole_frac_vap_sat, self.enth_mol_vap, self.entr_mol_vap, self.vol_mol_vap]
62.     script_dir = os.path.dirname(__file__)
63.     self.pysmo_surrogate = PysmoSurrogate.load_from_file(
64.         os.path.join(script_dir, "rbf_HA_10000.json")
65.     )
66.     self.surrogate = SurrogateBlock()
67.     self.surrogate.build_model(
68.         self.pysmo_surrogate,
69.         input_vars=inputs,
70.         output_vars=outputs,
71.     )
72.
73.     self.water_props = HelmholtzParameterBlock(
74.         pure_component="h2o", phase_presentation=helmPhaseType.L, state_vars=StateVars.TPX
75.     )
76.     # call the methods that have to be constructed automatically.
77.     self._enth_mol()
78.     self._entr_mol()
~~~~~
2.     def _entr_mol_phase(self):
3.         def _rule_entr_mol_phase(b, p):
4.             if p == "Liq":
5.                 return b.water_props.stpx(T=b.temperature, p = b.pressure) #From Helmholtz EoS
6.             elif p == "Vap":
7.                 return b.entr_mol_vap #From the Surrogate
8.
9.         self.entr_mol_phase = Expression(
10.            self.params.phase_list,
11.            rule=_rule_entr_mol_phase,
12.        )
#Equation converting mols to mass
```

```

14.     def _entr_mass_phase(self):
15.         def _rule_entr_mass_phase(b, p):
16.             if p == "Liq":
17.                 return b.enth_mass_phase[p] * b.params.mw_comp["water"]
18.             elif p == "Vap":
19.                 return b.entr_mol_vap * b.ave_MW_vap
20.
21.
22.         self.entr_mass_phase = Expression(
23.             self.params.phase_list,
24.             rule=_rule_entr_mass_phase,
25.         )

```

Figure 38 Code excerpt showing the injection of the surrogate and Helmholtz model.

5.4.8 Discussion

The final surrogate model demonstrated excellent agreement with established thermodynamic models, achieving over 99% alignment with both the CoolProp model for humid air and the CoolProp formulation for the liquid water phase. The combined use of dimensionality reduction, adaptive sampling, and Kriging-based surrogate modelling proved to be highly effective, delivering accurate performance across the entire solution domain (except for regions exhibiting discontinuous behaviour). Overall, the methodology significantly reduced the complexity and specialist knowledge typically required to implement detailed thermodynamic property models in equation-oriented frameworks.

While the approach is, in principle, is broadly applicable to a wide range of thermodynamic systems, its implementation remains somewhat labour-intensive. This is due to the continued need for users to manually define and code many routine process engineering relationships, such as material balances, phase conversions, and basic thermodynamic identities. However, these auxiliary calculations are often repeatable and thus lend themselves well to abstraction. Incorporating such recurring relationships into reusable, callable functions would streamline the implementation. For example, by automatically handling unit conversions or molar-to-mass transformations during property evaluations since in Figure 38. This would enable engineers to focus on high-level model configuration rather than repetitive low-level coding, thereby making the method more accessible and scalable for wider industrial adoption.

Another consideration is data availability. Whilst a significant number of thermodynamic studies are available for food and other non-petrochemical compounds, these studies often are not comprehensive. Milk for example is treated at conditions ranging from 4°C to 220°C and has significant variances in composition for site to site. Therefore, creating a surrogate from clean model data would not be applicable. Instead using the recorded process conditions to train the model would be required. This would add additional challenges due to the non-clean and often spotty nature of the data thus requiring additional techniques such as automated data cleaning and physics informed machine learning) to be implemented to manage this complication.

Another important consideration in the development of surrogate thermodynamic models is the availability and quality of data. While a substantial body of thermophysical research exists for food-related and non-petrochemical compounds, such datasets are often incomplete or limited in their operational scope. For example, milk is processed under a wide range of thermal conditions, from refrigeration at approximately 4 °C to drying above 220 °C. Furthermore, its composition varies significantly between processing sites due to differences in raw milk supply, formulation, and product requirements. These factors render the use of clean, laboratory-based datasets for surrogate training largely impractical. Instead, the surrogate models must be trained using historical or operational process data captured from real plant environments.

Real-world process data also presents its own set of challenges. Such data are often noisy, incomplete, or inconsistently logged, and may contain measurement artefacts or anomalies. As a result, robust data preprocessing becomes essential. This includes the use of automated data cleaning techniques, outlier detection, and interpolation methods to ensure the training data are suitable for model development. In addition, the integration of physics-informed machine learning offers a promising approach to mitigate some of the issues inherent in sparse or imperfect datasets by embedding known physical laws directly into the surrogate model structure. This hybrid strategy helps maintain consistency with first-principles thermodynamics while leveraging the flexibility of data-driven modelling to represent complex, site-specific behaviours.

5.5 Conclusion

This chapter has addressed one of the most fundamental and technically challenging enablers of PI&E analysis: the creation of accurate, solver-compatible thermophysical property packages. Building on the digital twin framework established in Chapter 4, the work presented here has developed and validated property models spanning three complementary approaches activity coefficient models for

complex food fluids, Helmholtz-based formulations for high-accuracy pure substances, and machine-learned surrogates for computationally complex or compositionally variable mixtures.

The first strand of work demonstrated how conventional activity coefficient models, implemented with carefully defined pseudo-components, can be adapted to represent multicomponent dairy fluids such as milk. This required the integration of sparse and heterogeneous literature data into a coherent, internally consistent formulation, with validation against industrial operating conditions and measured thermal performance. The methodology highlights both the feasibility and the limitations of extending petrochemically-oriented commercial simulators into the food sector, particularly in the absence of comprehensive vapour–liquid equilibrium data.

The second strand focused on the implementation of high-fidelity Helmholtz-energy equations of state for pure fluids of strategic relevance to electrification technologies, notably ammonia and natural refrigerants. By embedding these models into the open-source, equation-oriented IDAES framework, it was possible to achieve near-reference accuracy across the full thermodynamic spectrum, while maintaining solver-friendliness through smooth, differentiable formulations. Validation against NIST reference data confirmed their suitability for optimisation-driven simulation and design.

The third strand addressed the persistent challenge of scalability and accessibility by developing a generalisable workflow for surrogate property modelling. Using humid air as an industrially relevant case study, a hybrid approach was implemented that combined dimensionality reduction, adaptive sampling, and Kriging-based regression to capture complex, non-ideal behaviour within a solver-compatible, data-driven model. This approach offers a practical pathway for process engineers to generate custom property packages from plant-specific or operational data, thereby extending the applicability of PI&E to fluids and mixtures for which no high-fidelity analytical model exists.

Collectively, these developments close a critical capability gap identified in Chapter 4: the absence of property packages capable of accurately and efficiently representing the fluids encountered in low- to medium-temperature process heat applications, particularly in the food and dairy sector. The outputs of this chapter are twofold:

1. A suite of validated, high-fidelity property packages integrated directly into the digital twin, enabling thermodynamically robust targeting, design, and evaluation studies; and
2. A reproducible methodology for property model development that balances accuracy, solver compatibility, and user accessibility.

With these property packages in place, the digital twin described in Chapter 4 is now both structurally complete and thermodynamically rigorous. The subsequent chapters will leverage this enhanced simulation environment to conduct advanced PI&E analyses, quantify electrification opportunities, and evaluate process retrofit and optimisation scenarios with a high degree of confidence.

Chapter 6 Digital twin

6.1 Introduction

The development of a design digital twin represents the transition from preparation and model definition into a fully operational simulation environment capable of supporting electrification and integration studies. In the preceding chapters, the intent of the digital twin was established, system boundaries were defined, and advanced thermophysical property packages were developed to accurately represent complex fluids such as milk and humid air. These elements form the essential foundation for building a digital twin that is both technically rigorous and practically applicable. The present chapter advances this foundation by assembling the complete process model, embedding the unit operations, and establishing the solving structures necessary to enable analysis of an electrified milk evaporation system.

The creation of the digital twin requires careful translation of industrial process features into simulation platforms. On one hand, commercial tools such as Aspen HYSYS and DWSim provide accessible environments for flowsheet construction and rapid prototyping of unit operations. On the other, open-source equation-oriented environments, such as the Ahuora platform powered by IDAES, offer the flexibility to define custom unit models, incorporate novel property formulations, and integrate directly with optimisation workflows. Employing both classes of platforms highlights the complementarity between ease of use and analytical depth, while also demonstrating the platform-agnostic nature of the Process Integration and Electrification (PI&E) methodology.

This chapter begins by introducing the modelling of critical unit operations, including ejectors, direct steam injection (DSI), direct contact heating (DCH), mechanical vapour recompression (MVR), falling-film evaporation, and high-temperature heat pumps. Each unit operation must be represented not only in terms of steady-state performance but also in a form that preserves thermodynamic consistency across the integrated system. Particular emphasis is placed on ensuring that energy flows and exergy relationships are transparent and traceable (an essential requirement for subsequent targeting and optimisation).

Following the construction of unit models, attention is given to assembling the wider flowsheet and defining the solving structures. This involves selecting between sequential-modular approaches, as seen in HYSYS, and equation-oriented approaches, as implemented in IDAES. Each method presents advantages and limitations in terms of convergence behaviour, solver stability, and flexibility for

integration analyses. The discussion therefore explores not only the technical steps of building the flowsheet but also the methodological trade-offs inherent in different simulation paradigms.

Finally, the chapter concludes by demonstrating the completion of the digital twin across both environments, enabling a comparative evaluation of their capabilities. The outcome is a functional, validated digital twin of the milk evaporation process, constructed in HYSYS, DW Sim and replicated within the Ahuora platform powered by IDAES.

While the digital twin is later used as the foundation for PI&E analysis, it is also a stand-alone modelling artefact. Once constructed, the model can be used independently to investigate process behaviour, compare simulation platforms, test alternative operating conditions, and extract structured process data for other engineering analyses. The PI&E studies presented in Chapters 7 and 8 therefore represent one application of the digital twin rather than its only purpose.

6.2 Methods for Creating and Evaluating Unit Models

Most commercial and open-source process simulation frameworks such as Aspen Plus, HYSYS, and IDAES offer extensive libraries of standard unit operations including pumps, heat exchangers, flash drums, and distillation columns. However, in specialised applications such as dairy evaporation or process heat recovery, certain unit operations may be absent or insufficiently detailed. A key example is the thermo-vapour recompression (TVR) or ejector, which is commonly used in multi-effect evaporation systems to improve steam economy. Despite its widespread industrial use, TVR ejectors are often omitted from default unit operation libraries. The following section detail how the TVR unit model was generated.

6.2.1 Ejector / TVR modelling method

The key literature used to develop the TVR system model comes from Friso (2022). Some adaptations were made to the method to integrate the IAPWS Helmholtz property models rather than ideal gas formulations previously used.




Figure 39 The general schematic of a TVR system sourced from Friso (2022).

A thermo-vapour recompression (TVR) ejector is a device used to increase the pressure of a low-pressure vapour stream, known as the induced fluid (denoted by subscript i), by utilising the energy from a higher-pressure motive fluid (denoted by subscript m). In the context of dairy processing, the induced fluid typically consists of low-pressure water vapour generated during evaporation, while the motive fluid is intermediate-pressure steam, generally in the range of 8–12 bar. The TVR functions by directing the motive steam through a converging-diverging nozzle, causing it to expand and accelerate to supersonic velocities. This creates a vacuum that entrains the induced vapour. The two fluid streams then mix in a constant-area chamber before being decelerated in a diffuser, resulting in a recompressed vapour stream at an intermediate pressure. The fundamental governing principle for this mixing process is the conservation of momentum. Neglecting friction and pressure losses, the momentum balance across the ejector stages can be expressed as shown in Equation 32. In this formulation, variables denoted by subscript 1 refer to the motive fluid as it exits the nozzle; subscript 2 refers to the state at the entry of the mixing chamber; subscript 3 represents conditions at the exit of the mixing section; and subscript 4 denotes the state at the diffuser outlet.

$$\omega_1 G_m + \omega_2 G_i = \omega_3 (G_m + G_i) \quad \text{Equation 32}$$

By applying the conservation of momentum, the entrainment ratio can also be expressed in terms of the fluid velocities. To relate these velocities to thermodynamic properties, the principle of energy conservation is applied. Specifically, the kinetic energy of the fluid streams is derived from their isentropic enthalpy change during expansion or acceleration. This yields a series of expressions, Equation 33, Equation 34 and Equation 35.

$$\omega = \frac{G_i}{G_m} = \frac{\omega_1 - \omega_2}{\omega_3 - \omega_2} \quad \text{Equation 33}$$

$$\omega_1 = \sqrt{2 \eta_{E1} \Delta h_1} \quad \text{Equation 34}$$

$$\omega_2 = \sqrt{2 \eta_{E2} \Delta h_2} \quad \text{Equation 35}$$

$$\omega_3 = \sqrt{2 \frac{\Delta h_3}{\eta_D}} \quad \text{Equation 36}$$

The modelling approach presented here deviates from the method outlined by Friso (2022). In Friso's previous work, the ejector was modelled under the assumption of ideal gas behaviour, with iterative calculations performed over a range for P_2 to determine the maximum achievable entrainment ratio. In contrast, the objective within the PI&E framework is to estimate the quantity of intermediate-pressure (IP) utility steam required to recompress the low-pressure vapour stream thus characterising the thermo-vapour recompression (TVR) device in terms of utility integration.

Lie (2014) conducted a comprehensive review of ejector models and correlations available in the literature. Their findings suggest a typical isentropic efficiency of approximately 0.7 for the suction and mixing sections of steam ejectors. This is in contrast to the efficiency assumptions made by Friso (2022), which ranged from 0.92 to 0.99. The discrepancy is likely attributed to the significant pressure differential encountered in this industrial application, where motive steam at 12 bar is used to compress vapour at vacuum conditions (e.g., 0.112 bar).

To accommodate these practical conditions, the present approach assumes a suction pressure drop of 0.035 bar and applies isentropic expansion relations to determine the thermodynamic properties of the involved streams. The diffuser section efficiency is estimated based on the Mach number of the expanded motive stream, as illustrated in Figure 40. The Mach number, defined as the ratio of fluid velocity to the local speed of sound, which the latter is calculated using the IAPWS-95 formulation for water. The velocity of the fluid is derived from the isentropic enthalpy change, as outlined in Equation 34.




Figure 40 Experimental relationship between the diffuser efficiency and Mach number of the expanded motive gas sourced from Friso (2022).

The next section to model is the mixing section. Again, the conservation of energy principle can be used to mix the induced and motive streams states. It is also important to account for the enthalpy energy that was converted to kinetic energy during expansion as this would not contribute to the total mixed enthalpy as shown in Equation 37. An initial entrainment ratio is used to initialise the mass and energy balance for solving.

$$h_{mix} = \frac{h_{i,2}G_i + h_{m,2}G_m}{G_i + G_m} - \frac{1}{2} \left[\frac{v_{2m}}{\left(1 + \frac{G_i}{G_m}\right)} \right]^2 \quad \text{Equation 37}$$

To model the diffuser section first the user inputs a target outlet pressure P_4 . Subsequently the isentropic compression formula is used determine the final outlet state of the diffuser and therefore the full TVR unit. The final part of the model is to solve for the entrainment ratio that satisfies the conservation of momentum principles that are detailed above as well has the overall energy balance of the system $H_{out} = H_i + H_m$.

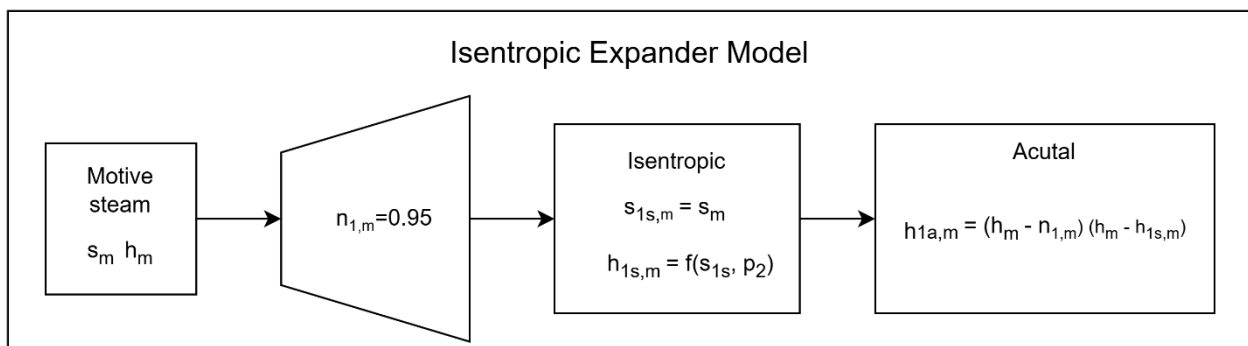
An example and verification calculation are provided to show how the method can be calculated. Friso (2022) gathered many ejector studies across a range of working fluids. One of these studies is for a low-pressure steam ejector at the following boundary conditions. Motive steam at 2.7 bar and 130°C, induced vapour at 0.012 bar 10°C and target pressure of 0.05 bar. The efficiencies used for

the motive steam and induced expansion were 0.95 and 0.99 respectively based on Friso's (2022) data.

Step one: calculate thermodynamic states using IAPWS Helmholtz formulation. Assume the entrainment ratio is one to initialise.

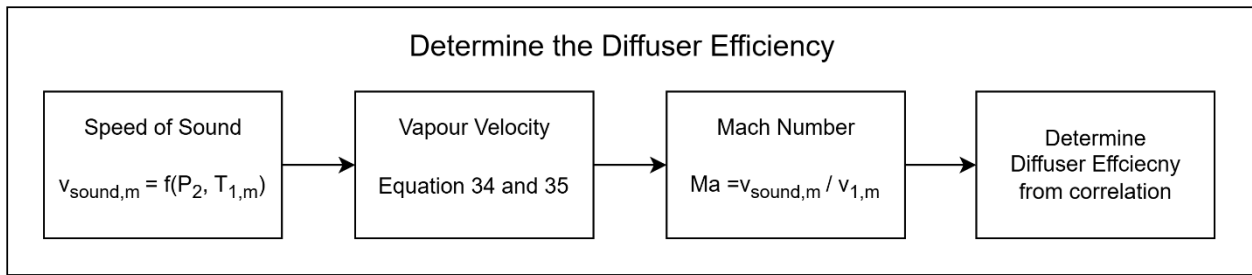
Motive in			Induced in		
G_m	1	kg/s	G_i	1	kg/s
P_m	2.7	bar	P_i	0.0123	bar
T_m	130	C	T_i	10	C
h_m	2722	kJ/kg	h_i	42	kJ/kg
s_m	7.03	kJ/kgK	s_i	8.9	kJ/kgK
H_m	2722	kW	H_i	42	kW

Step two: Apply isentropic expansion model for both the motive and induced streams using 99% of P_i as the representative pressure drop for P_1 which is equivalent to P_2 . $\eta_{1,m} = 0.95, \eta_{1,i} = 0.99, P_{1 \text{ and } 2} = 0.0122$



Expanded Motive			Expanded Induced		
$G_{1,m}$	1	kg/s	G_{1i}	1	kg/s
$P_{1,m}$	0.0122	bar	P_{1i}	0.0122	bar
$T_{1,m}$	9.85	C	T_{1i}	9.85	C
$h_{1,m}$	2026	kJ/kg	h_{1i}	2493	kJ/kg
$s_{1,m}$	7.16	kJ/kgK	s_{1i}	8.9	kJ/kgK
$H_{1,m}$	2026	kW	H_{1i}	2493	kW

Step three: Calculate the Mach number and determine a corresponding diffuser efficiency. Use Helmholtz package to calculate speed of sound



$$v_{1,m} = \sqrt{2(2722 \text{ kJ/kg} - 2026 \text{ kJ/kg}) * 1000} \Rightarrow 1180 \text{ m/s}$$

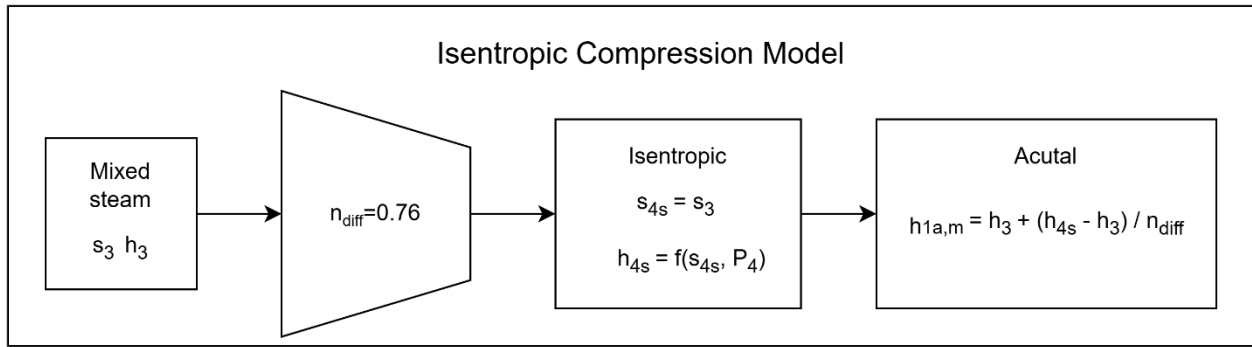
$$v_{1,sound} = f_{helm}(0.012 \text{ bar}, 9.85 \text{ }^\circ\text{C}) \Rightarrow 416.8 \text{ m/s}$$

$$Ma = \frac{1180 \text{ m/s}}{416.8 \text{ m/s}} \Rightarrow 2.8 \Rightarrow \text{Correlation} \Rightarrow \eta_{diff} = 0.76$$

Step four: Perform mixing utilising Equation 37.

Mixing Chamber		
G ₃	2	kg/s
P ₃	0.0122	bar
T ₃	9.85	C
h ₃	2086	kJ/kg
s ₃	7.37	kJ/kgK
H ₃	4171	kW

Step five: Apply isentropic compression model for the mix stream using the calculated. P₄ (target pressure) = 0.05 bar



Diffuser Exit		
G ₄	2	kg/s
P ₄	0.05	bar
T ₄	32.9	C
h ₄	2299	kJ/kg
s ₄	7.54	kJ/kgK
H ₄	4599	kW

Step six: Create a constraint to enforce the mass and energy balance of the system and solve for entrainment ratio

$$H_4 = H_m + H_i \Rightarrow H_m + H_i = 2764 \text{ kW}$$

$$4599 \text{ kW} \neq 2764 \text{ kW} \Rightarrow \omega \text{ is not correct}$$

$$\text{Solve for } \omega \Rightarrow \omega = 0.386$$

Friso (2022) quotes that the experimental entrainment ratio for the device was 0.4. With the current assumptions the proposed mathematical model estimated an entrainment ratio of 0.386 which is a 3.56 % difference. The method was further validated using industrial data from a real milk evaporation system. The device operated with a motive IP steam stream of 12 bar saturated and an induced vapour stream of 0.122 bar saturated. The target pressure for the TVR was 0.19 bar. The efficiency for the suction section was assumed to be 0.7 as suggested by Lie (2014). The model estimated the entrainment ratio to be 3.33 with the industrial data measuring an average entrainment ratio of 3.25 which is a 2.43% difference.

This accuracy was highly encouraging and good enough to proceed with the PI&E method. One should note that the users' assumptions for the isentropic efficiencies will have an impact on the accuracy of the model. These efficiencies will be subject to the device's geometry and boundary conditions for the device.

6.2.2 Direct Steam Injector (DSI)

A DSI is a heat transfer unit operation commonly employed in the dairy industry to achieve rapid and uniform heating of milk. In this system, steam is injected directly into the milk stream, where it condenses and transfers its latent heat to the fluid, enabling extremely fast temperature increases. This makes DSI particularly well suited for high-temperature applications such as Ultra-High Temperature (UHT) processing and pasteurisation.

The modelling approach for a DSI is relatively straightforward, involving a combined mass and energy balance.

The modelling of the DSI unit operation is based on the assumption of:

- Instantaneous and complete condensation of steam.
- Adiabatic operation (no heat loss to the surroundings).
- Perfect mixing of steam and milk.
- The outlet pressure is set to be the minimum value of the inlet pressures

Let the milk stream be denoted by subscript m, the steam stream by s, and the mixed outlet stream by o.

The total mass flow rate of the outlet stream is the sum of the inlet mass flow rates:

$$\dot{m}_o = \dot{m}_m + \dot{m}_s \quad \text{Equation 38}$$

For each component i (e.g., water, lactose, protein, fat):

$$\dot{m}_o x_{\{i,o\}} = \dot{m}_m x_{\{i,m\}} + \dot{m}_s x_{\{i,s\}} \quad \text{Equation 39}$$

Assuming adiabatic operation, the enthalpy of the outlet stream equals the combined enthalpy of the inlet streams:

$$\dot{m}_o h_o = \dot{m}_m h_m + \dot{m}_s h_s \quad \text{Equation 40}$$

Here:

- h_m : specific enthalpy of the milk stream (J/kg)
- h_s : specific enthalpy of the steam stream (J/kg) — includes both sensible and latent heat
- h_o : specific enthalpy of the outlet mixture (J/kg)

However, a complication arises from the mixing of two distinct fluids: pure steam, modelled using a Helmholtz-based equation of state, and milk, a multi-component fluid modelled using the activity coefficient model NRTL. Although water is a component in both streams (allowing the component balance to be handled directly) the enthalpy calculation requires careful consideration. Specifically, the reference state used in each property model must be aligned to avoid introducing artificial offsets in enthalpy values. In this work, most property packages were standardised to use the triple point of water as the enthalpy reference state, thereby ensuring thermodynamic consistency. If models with differing reference states were used, appropriate adjustments or translations would be required.

A specific challenge encountered during property model integration was the misalignment of reference enthalpies between the milk formulation and the IAPWS formulation used for pure water. The heat capacity (C_p) and corresponding enthalpy correlations for milk components (where enthalpy is the integral of C_p) were developed using a reference state of 25°C and 1 atm for the gas phase. In contrast, the Helmholtz-based formulation used for water adopts the triple point of water as its reference state.

To address this discrepancy, a comparison tool was incorporated alongside the DSI unit operation model. This tool checks for enthalpy deviations by comparing common fluid components across both property packages. As illustrated in Figure 41, the liquid-phase enthalpy values for water in the milk formulation and the Helmholtz model were aligned to the same reference state (noting that the milk stream contains negligible milk solids in the liquid phase). However, a significant deviation was observed in the vapour phase, highlighting the importance of reference state consistency when combining fluids across different thermodynamic models.




Figure 41 Enthalpy comparison between milk and water showing deviation

To resolve the enthalpy misalignment, a reference enthalpy offset term (denoted as H in Figure 42) was adjusted until the enthalpy values of the two fluids were brought into alignment, as demonstrated in Figure 43. The only remaining discrepancy occurs in the liquid–vapour transition region, where some smoothing is observed. This is expected, given that the milk property package was designed to model multicomponent mixtures, whereas the Helmholtz formulation targets pure components.

While the approach adopted here may be considered a simplification, it performs adequately for the specific use case presented. In more complex scenarios involving additional components a more rigorous correction method or the intervention of an experienced user may be necessary. However, by embedding this adjustment mechanism into a digital platform, the procedure can be standardised. This enables users to confidently select and apply suitable property packages without requiring a large degree of expertise in thermodynamic references.

```

1.         "cp_mol_ig_comp_coeff": {
2.             ...
23.         "F": (-250.8810, pyunits.kJ / pyunits.mol),
24.         "G": (202.3, pyunits.J / pyunits.mol / pyunits.K),
25.         "H": (-241.83, pyunits.kJ / pyunits.mol),
26.     
```

Figure 42 Code snippet showing the enthalpy offset term




Figure 43 Enthalpy comparison between milk and water showing deviation

6.2.3 Direct Contact Heater (DCH)

As part of the milk heat treatment section, the evaporation system incorporates a DCH. In this unit, preheated milk out of the direct steam injector at 95 °C is introduced into a chamber maintained below its saturation pressure, causing partial flashing. The resulting vapour rises and directly contacts cooler incoming milk sprayed from the top of the chamber, facilitating efficient heat transfer. The technological benefit of direct contact heating is its relatively fast heating rates, due to the nature of the hot fluid directly combining with the cold fluid. The faster heating rate therefore means less time for bacterial growth in the equipment and subsequently more useful time in operation.

The key to accurately modelling the DCH lies in the stable and precise simulation of the milk's bubble point. A reliable two-phase model enables the calculation of both the enthalpy contributions from the vapour and liquid phases, as well as the proportion of each phase present. Once the flash behaviour is captured accurately, the subsequent modelling steps involve phase separation and isobaric mixing, following a similar approach to that used for the DSI.


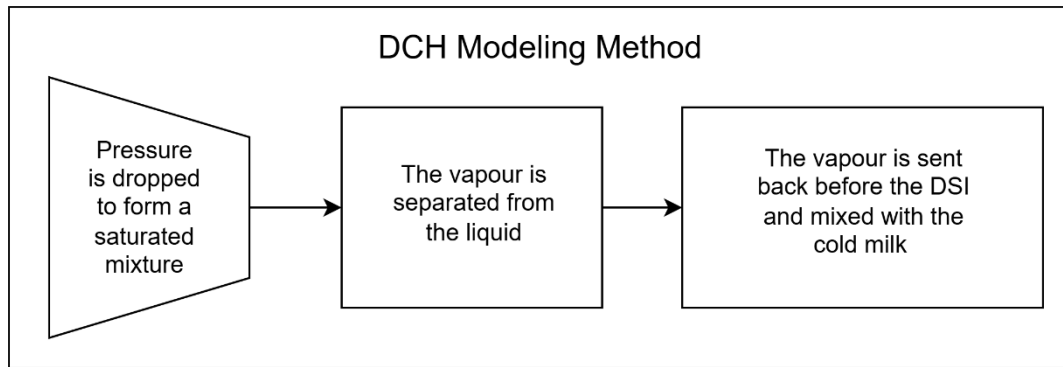


Figure 44 DCH Model

6.2.4 Mechanical Vapour Recompressor (MVR)

In an MVR system, a low-pressure vapour typically generated from a boiling liquid is compressed using a mechanical compressor to a higher pressure and temperature. In the context of an evaporator this recompressed vapour can then be reused as a heating medium, replacing or significantly reducing the need for fresh utility steam (like the TVR). This process enables the reuse of the vapour's latent heat, which would otherwise be wasted, by upgrading it to a more useful energy state.

Mechanical Vapour Recompression (MVR) operates on the principle of isentropic compression. Where $h_{2,s}$ is the enthalpy calculated by the property package evaluated at the outlet pressure P_2 and the entropy on the inlet s_1 .

$$\eta_s = \frac{h_{2,s} - h_1}{h_{2,a} - h_1} \quad \text{Equation 41}$$

$$w_s = h_2 - h_1 \quad \text{Equation 42}$$

The isentropic efficiency of a mechanical vapour recompression (MVR) unit can, in principle, be evaluated by embedding a performance correlation (typically a function of the pressure ratio). However, analysis of industrial data indicates that the isentropic efficiency remains relatively stable for well-sized systems, with a typical value around 80%.

An ancillary component commonly included in MVR systems is the *wetting nozzle*, which serves both operational and thermal roles. This nozzle introduces a small quantity of liquid water directly onto the compressor blades to prevent fouling from residual milk solids. The injected water mixes with the vapour, and this mixing (modelled using standard mixing equations) has a cooling effect by lowering the specific enthalpy of the vapour stream. Therefore, since the vapour is generally superheated after compression, the injection of cooler liquid water provides a secondary superheating effect that enhances downstream heat transfer performance.




Figure 45 Representation of the wetting nozzle air represented by large dark arrows with water being sprayed

6.2.5 Falling Film Evaporator

A falling film evaporator is a widely utilised thermal separation unit in the dairy, food, and chemical processing industries, particularly where the preservation of heat-sensitive product quality is essential. The fundamental operating principle involves distributing the liquid feed evenly along the top of a vertically oriented bundle of heated tubes. The liquid flows downward as a thin film along the inner surface of the tubes, driven by gravity.

Heat is applied to the outer tube surfaces causing a portion of the liquid film to evaporate as it descends. The generated vapour flows co-currently with the liquid film toward the evaporator's lower outlet to the vapour-liquid separator, typically positioned downstream of the tube bundle. As the mixture exits the bottom of the evaporator tubes, it enters a larger chamber where the reduction in velocity facilitates the gravitational separation of the heavier liquid phase from the lighter vapour phase.

The falling film evaporator is implemented as a two-sided heat exchanger model. The formulation adopts a steady-state approach using the $Q = U A \Delta T_{LM}$ method, which is well suited to simulation of a heat exchange process. In this case, heat is transferred from a recompressed vapour stream (hot side) to a downward-flowing milk film (cold side), leading to partial evaporation of water from the milk.




Figure 46 Schematic of the inlets and outlets for a milk evaporator

The model defines two independent control volumes. The hot side, representing the recompressed vapour stream (e.g., from an MVR or TVR) and the cold side, representing the milk undergoing evaporation. Each control volume supports a full separate thermodynamic property calculation through the specified property package tailored for the fluid. This presents a challenge as the Helmholtz package is better suited to the hot side (a better characterisation of entropy for isentropic calculations embedded in the TVR and MVR models), but the vapour is being recycled from the activity coefficient model for milk.

Therefore, an intermediate translation step is required to convert between the two models. The simplest translation is constraining the pressure and temperature calculated by the milk package to equal the pressure and temperature being inputted into the water package. Since the recycled vapour stream exists only as a vapour (not in the saturated zone) the other properties can be calculated simply.

The core formulation defines key variables including the overall heat transfer coefficient (U , in $W/m^2 \cdot K$), the effective heat transfer area (A , in m^2), and the temperature differences at the inlet and

outlet ends of the exchanger (ΔT_{in} , ΔT_{out}). These temperature differences are used to calculate the log-mean temperature difference (ΔT_{LM}) through a user-defined callback function, where a user can describe if the flow is co-current, counter-current or cross flow.

To ensure energy conservation, the model imposes a unit-level energy balance constraint that enforces equivalence between the heat lost by the hot side and the heat gained by the cold side (Equation 43).

$$Q_{hot} + Q_{cold} = 0 \quad \text{Equation 43}$$

This heat-exchanger model is utilized for other heat exchange operations as well.

6.2.6 Heat Pump modelling

Heat pumps are central to electrification strategies in industrial process systems, offering the ability to recover and upgrade low-grade thermal energy to useful temperature levels using electricity as the primary energy input. Their integration into process models requires accurate characterisation of both their thermodynamic behaviour and their interaction with surrounding unit operations. In this research, two complementary approaches are adopted for heat pump modelling: rigorous cycle simulation and performance-curve-based models.

The first approach involves rigorous simulation of the thermodynamic cycle. This method allows a more detailed evaluation of different working fluids and cycle configurations. It accounts for non-ideal thermodynamic behaviour, potential pressure drops, isentropic efficiencies, superheating and subcooling, and real fluid properties using the Helmholtz-based equations of state. This level of detail is necessary when deciding what type of cycle will be most performant for the chosen application (i.e. the selection between an ammonia cycle or a propane cycle with a sub cooler).

The most complex of heat pump cycles investigated, was a single-stage vapour compression cycle incorporating an internal heat exchanger (IHX) and sub-cooler.

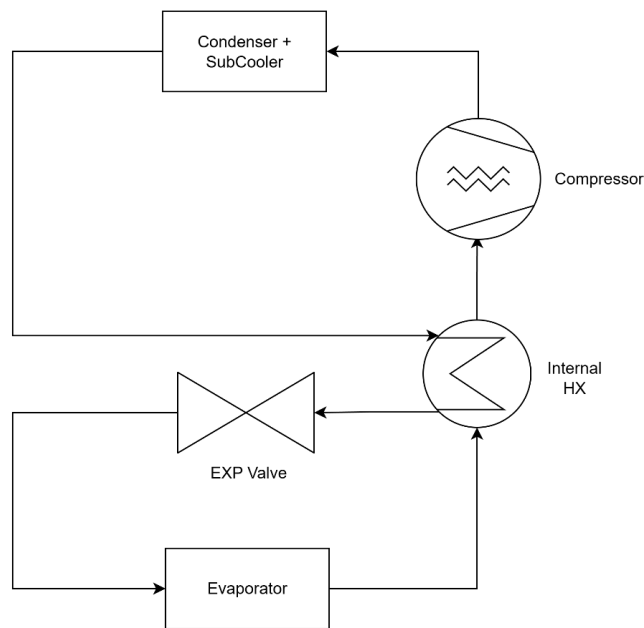


Figure 47 Process flow diagram of an internal heat exchange with sub cooling cycle

The cycle includes the following key components:

- Evaporator: where heat is absorbed from the process or waste stream (usually ambient air) (see 6.2.5 Heat Exchanger Model)
- Compressor: raises the vapour pressure and temperature (see 6.2.4 Compressor Model)
- Condenser and sub cooler rejects heat to the sink (which is usually utility hot water) (see 6.2.5 Heat Exchanger Model plus 1D additions described below)
- Expansion valve: reduces pressure and initiates partial vaporisation. This is modelled using the adiabatic assumption $\dot{m}_i h_i = \dot{m}_o h_o$. A new outlet pressure is specified and this new pressure alongside the same specific enthalpy is used to calculate the rest of the system.
- Internal heat exchanger: transfers heat from the hot liquid (post-condenser) to the cold vapour (post-evaporator) prior to compression and expansion (see 6.2.5 Heat Exchanger Model)

One of the key challenges in simulating heat pump cycles is accurately capturing the non-linear heat transfer behaviour of the refrigerant and utility streams within the heat exchangers. This is particularly evident with fluids such as CO₂, which exhibit pronounced thermophysical property changes near the critical region. As the refrigerant approaches its critical point, the temperature–enthalpy profile becomes highly non-linear, leading to a non-uniform temperature driving force along the length of

the heat exchanger. As shown in Figure 48, this can result in the temperature pinch point (the location of minimum temperature difference between the hot and cold streams) shifting away from the inlet or outlet and occurring instead within the interior of the exchanger. This behaviour complicates both the design and numerical stability of the model, as conventional assumptions of monotonic temperature gradients or fixed pinch locations may no longer hold. This nonlinear behaviour also occurs when fluids undergo phase change.

To address the challenges associated with non-linear heat transfer profiles a one-dimensional (1-D) heat exchanger model was employed. Unlike lumped-parameter models that assume uniform temperature driving forces, a 1-D heat exchanger discretises the heat exchanger length into a series of finite control volumes, allowing temperature, pressure, enthalpy, and phase behaviour to vary along the flow path. This approach enables accurate tracking of the local temperature difference between the hot and cold streams at each segment, which is essential for capturing shifting pinch points. The model solves coupled mass, energy, and momentum balances across each segment, providing a high-fidelity representation of the thermal interaction between streams.

Like the 0D heat exchanger model two separate Helmholtz property packages are defined, one for the refrigerant and one for the process fluid. The heat exchanger is defined with counter-current flow pattern.

To represent spatial variation along the exchanger, a continuous set is introduced to define the characteristic length of the control volume. This domain is normalised between 0 and 1, creating a non-dimensional spatial variable $x \in [0,1]$ along the exchanger length. This continuous set forms the basis for defining the spatial distribution of temperature, pressure, enthalpy, and other properties throughout the exchanger.

The governing equations (mass and energy balances, and the heat transfer performance equation) are conceptually similar to the 0D model but must now be expressed as functions of the continuous spatial coordinate. The standard heat transfer relation is modified to account for local variation in temperature and heat flux:

$$Q_{hot,x} = -U_x \frac{A}{L_{hot}} (T_{hot,x} - T_{cold,x}) \quad \text{Equation 44}$$

where x is the non-dimensional position along the exchanger, U_x is the local heat transfer coefficient, and L_{hot} is the total length of the hot side. This formulation enables local temperature differences and gradients to be resolved throughout the exchanger.




Figure 48 A CO₂ cycle exhibiting multiple pinch points throughout the length of the heat exchanger

To enable numerical solution of the 1D heat exchanger model, the spatially continuous domain must first be discretised. This is necessary because standard nonlinear optimisation solvers operate on algebraic systems, not on models defined with continuous sets and derivatives. Pyomo provides two primary approaches for performing this transformation: finite difference methods and collocation schemes. Both methods convert the differential-algebraic equations (DAEs) defined over a ContinuousSet into a system of algebraic equations by introducing a discrete set of points along the domain.

Pyomo supports two main families of discretisation methods for transforming models with continuous domains (Nicholson et al., 2018): finite difference methods and collocation schemes. Both approaches aim to convert differential-algebraic models into algebraic systems that can be handled by standard nonlinear solvers, but they differ in accuracy, complexity, and computational efficiency.

Finite Difference Methods (FDM): These are conceptually simple and computationally efficient. Derivatives are approximated using local differences between adjacent discretisation points. While easy to implement and suitable for many engineering applications, finite difference schemes may require a large number of elements to achieve sufficient accuracy (particularly in systems with steep gradients or non-linear property variations). In such cases, numerical diffusion and reduced resolution of critical phenomena (e.g., pinch points) can be a limitation.

Collocation Methods: Collocation techniques approximate the solution within each finite element using polynomial basis functions and enforce the governing equations at specific collocation points. The Lagrange-Radau collocation scheme, used in this work, offers higher-order accuracy with fewer elements and improved stability when solving stiff systems. This makes it particularly well suited to thermodynamically non-linear systems such as those encountered in the condenser, where accurate resolution of property gradients is essential. However, collocation methods are more computationally intensive and require more complex model transformations.

In this research, collocation was chosen to discretise the 1D heat exchanger model, as it provided the necessary balance of accuracy and solver robustness to capture the local thermal behaviour of fluids under subcritical, trans-critical, and supercritical conditions.

After trialling a variety of finite element discretisation's for this application it was determined that between 30 and 40 finite elements were sufficient to model the nonlinear behaviour. Also, the U was set a constant value of $3000 \text{ W/m}^2\text{K}$ across the exchanger as a simplifying assumption.

Once the unit models have been implemented a set of key degrees of freedom and performance constraints can be manipulated to solve the cycle.

Table 16 Variables to be changed by the user

Inputs	
Fluid	Comp η_{is}
Sink T outlet	Comp η_{mech}
Sink T inlet	ΔT_{min}
Fixed T evaporator	$\Delta T_{superheat,target}$
Degree of subcooling	$\Delta T_{sink,offset}$

This high-fidelity modelling approach is particularly valuable during the technology selection phase, where it enables rigorous evaluation of different refrigerants and cycle configurations to identify the most suitable heat pump system for process electrification. However, once a configuration has been selected, the detailed model (especially when coupled with multiple internal recycles) introduces unnecessary computational overhead and complexity within the broader evaporator flowsheet. This is true across both commercial and open-source modelling frameworks. To address this, it is proposed that the high-fidelity model be used only during configuration screening, after which it is replaced by a lightweight surrogate model. This surrogate can preserve the essential thermodynamic behaviour while significantly reducing solution time, thereby facilitating integration into larger-scale process simulations and optimisation studies.

To support efficient process-wide integration of the selected heat pump configuration, two surrogate modelling strategies were adopted. The first utilised PySMO (Lee et al., 2021), a surrogate modelling toolkit available in the IDAES framework, to train data-driven models (such as radial basis functions) on either high-fidelity simulation outputs (the model above) or experimental data, where available. This approach allows the surrogate to capture complex thermodynamic relationships over a defined operating window, enabling fast and solver-friendly evaluation of performance metrics such as heat duty, coefficient of performance (COP), and compressor power. A more comprehensive methodological description of the ML technique is given in (5.4.4)

The second approach involved the use of empirical performance correlations, which offer an alternative, pathway to estimate heat pump performance. These correlations relate COP and other key outputs to characteristic input parameters such as temperature lift and heat sink conditions, typically via simplified regression equations. These models are particularly well suited to light weight design simulation as they are representative of real-life equipment performance in a more generalised sense and where data is not available.

In this research, empirical models developed by Schlosser et al (2020a) were used as a practical and validated basis for performance estimation. Their study reviewed over 150 industrial heat pump installations and generated generalised COP correlations for multiple technology types, including standard subcritical vapour compression systems, trans-critical CO₂ cycles, and steam-generating heat pumps.

The COP is estimated using a four-parameter regression of the form:

$$\text{COP} = a \cdot (\Delta T_{\text{lift},m} + 2 \cdot b)^c \cdot (T_{h,\text{out}} + b)^d \quad \text{Equation 45}$$

where $T_{h,\text{out}}$, is the sink temperature leaving the condenser and $\Delta T_{\text{lift},m}$ is the mean temperature lift, calculated for subcritical systems as:

$$\Delta T_{\text{lift},m} = T_{h,\text{out}} - T_{l,\text{in}} \quad \text{Equation 46}$$

where $T_{l,\text{in}}$ is the source temperature entering the evaporator.

For trans-critical CO₂ heat pump cycles, conventional methods for estimating temperature lift and COP are often insufficient due to the complex thermodynamic behaviour of CO₂. Unlike subcritical systems, where the refrigerant undergoes phase change at nearly constant temperature and pressure, trans-critical cycles operate above the critical pressure, where no distinct liquid–vapour boundary exists. Instead, the refrigerant is cooled in a single-phase supercritical region, where temperature, density, and specific heat vary with pressure and enthalpy.

As a result, using a simple arithmetic mean temperature difference (as applied in subcritical systems) is inadequate, as it fails to represent the actual driving force for heat transfer. To address this, a more accurate metric based on the Lorenz cycle (a cycle that exhibits temperature glide) is used, which accounts for the changing specific heat and temperature profile of the supercritical CO₂ stream. Specifically, the logarithmic mean temperature is applied to represent the effective heat rejection temperature of the heat sink:

$$T_{h,m} = \frac{T_{h,\text{out}} - T_{h,\text{in}}}{\ln\left(\frac{T_{h,\text{out}}}{T_{h,\text{in}}}\right)} \quad \text{Equation 47}$$

This logarithmic mean temperature, $T_{h,m}$, serves as the average thermal level of the sink side and is then used to define the mean temperature lift:

$$\Delta T_{\text{lift},m} = T_{h,m} - T_{l,\text{in}} \quad \text{Equation 48}$$

These equations, when used with technology-specific fitting coefficients (as provided in), offer a computationally efficient means of evaluating COP based on a limited set of process variables.

Table 17 COP correlations for standard vapour compression, trans-critical CO₂ and steam generation heat pumps developed by Schlosser et al (2020a)

Heat Pump Type	Fitting Parameters (a, b, c, d)	Validity Range
SHP & HTHPs with R717 (ammonia)	a = 40.789, b = 1.0305, c = -1.0489, d = 0.29998	$0\text{ °C} \leq \theta_{l,in} < 40\text{ °C}$ $70\text{ °C} \leq \theta_{h,out} < 85\text{ °C}$ $30\text{ K} \leq \Delta T_{lift} \leq 75\text{ K}$
CO ₂ Heat Pump (transcritical)	a = 1261.2, b = 17.627, c = -1.3974, d = 0	$7\text{ K} \leq \Delta T_{lift} < 94\text{ K}$ $291\text{ K} \leq T_{h,m} < 343\text{ K}$ $0\text{ K} \leq \Delta T_{lift,m} \leq 74\text{ K}$

6.3 Methods for Deciding the Wider Solving Structures for the Total Flowsheet

Once property models and unit operations have been scoped and validated, the next critical step is solving the total integrated flowsheet. This involves linking all major unit operations such as the evaporator, heat exchangers, heat pump(s), compressors, and recycle loops into a coherent simulation structure.

The challenges associated with solving the full flowsheet depend heavily on the simulation platform used each presenting unique strengths and limitations.

6.3.1 Sequential modular simulation

Sequential modular simulation, used in tools like Aspen HYSYS and DWSim, breaks down the process flowsheet into individual unit operations, which are solved one at a time in a predefined sequence. The inputs to each unit are computed based on upstream units, and outputs are passed forward to downstream units. Recirculating streams (i.e., recycles) are handled by introducing tear streams, which are iteratively adjusted until convergence is achieved.

Two key challenges were faced when constructing using this methodology. The first is handling recycles and tight coupling. The flowsheet contains multiple recycle loops, including thermal and material recycles (e.g., vapour recompression, condensate return), which are strongly coupled and thermodynamically sensitive. In sequential modular frameworks, such loops can be numerically

fragile, particularly when fluid properties vary sharply with temperature or pressure which is common in multi-phase food fluids. Tear stream selection and convergence sequencing therefore become critical.

The primary strategy employed to improve convergence in the flowsheet was the careful selection of tear streams. Ideally, tear streams should be located in parts of the system where process variables (such as flowrate, temperature, and composition) are relatively insensitive to upstream changes. In contrast, streams located near phase boundaries, where even small perturbations can cause transitions between liquid and vapour, are poor candidates for tearing. In the evaporator system studied here, the presence of multiple internal recycles made it particularly challenging to identify suitable tear locations.

To manage this complexity, a stage-wise (or block-wise) solving approach was adopted. This method involves dividing the flowsheet into logical sub-processes (such as the evaporator, preheating, and COW recovery) and solving each independently under fixed boundary conditions before attempting to converge the full system. By doing so, the user can reduce the overall numerical complexity and propagate improved initial guesses from one block to the next. Importantly, this approach also provides natural boundaries for placing tear streams. Tears are most effective when located at interfaces between independently solvable subsystems, where local recycles have already been resolved and stream properties are stabilised. In such cases, the values used to initialise the tear streams are typically well-informed, improving the likelihood of successful convergence in the final system-wide solve.

Another major challenge encountered in the sequential modular framework was the limited flexibility for customisation, particularly when attempting to incorporate machine-learned (ML) surrogates, regressed performance curves, and non-standard unit operations. Aspen HYSYS and DWSim does not natively support direct integration of Python-based models or symbolic expressions, which are often required to implement data-driven correlations such as the empirical heat pump COP models developed from Schlosser et al (2020a). As a workaround, these models had to be embedded into the simulation via the inbuilt spreadsheet interface, resulting in a disjointed and error-prone workflow. This approach was similarly required for integrating cost curves and the custom ejector model, all of which demanded additional logic beyond what HYSYS and DWSim unit blocks could provide. The spreadsheet interface, while convenient for small parameter manipulations, becomes increasingly difficult to manage and scale when used to implement multiple equations or interconnected regressions with multiple variables from across the flowsheet.

6.3.1.1 Open-Sourced Python Based Equation Orientated Simulation

Unlike commercial sequential modular tools, IDAES requires the user to define every aspect of the model explicitly, including unit structure, state variables, constraints, and property formulations. This can be both powerful and burdensome: although it allows for simple integration of custom models such as the TVR, data-driven surrogates, and detailed thermodynamic packages, it also demands a deep understanding of both process systems engineering and nonlinear optimisation.

One key challenge is initialisation. Because the entire process model is solved simultaneously as a large nonlinear system, a poor initial guess or inconsistent variable bounds in even a single unit can prevent convergence. This necessitates again a block-wise build and solve strategy, where unit operations are first validated in isolation before being progressively assembled into the larger flowsheet. Additionally, the user must actively manage variable scaling, which is essential for solver stability, especially in systems with units of vastly different magnitude (e.g., pressure in MPa vs. enthalpy in J/mol). While IDAES provides utility functions for scaling and solver diagnostics, proper use requires iteration and user familiarity.

Another consideration is the steeper learning curve and increased development effort compared to commercial tools. Implementing a flowsheet in IDAES requires direct scripting in Python and familiarity with Pyomo, including concepts such as `Var`, `Constraint`, `Block`, and `ContinuousSet`. For example, to create the high-fidelity heat pump model, over 300 lines of code were required (this was for a single sub process).

The freedom to define every model component also means that the user must take full responsibility for ensuring that the system is well-posed. This requires careful planning and tracking of the degrees of freedom, particularly when combining multiple custom units, embedded performance equations, and internal recycles. Unlike commercial simulators that manage much of this implicitly, in IDAES every variable and constraint must be defined, fixed, or linked explicitly. If not managed properly, this can result in an under- or over-specified model, causing solver failure or incorrect results. The method to mitigate the specification issue is constructing visual representation with the planned fixed variables and initial values before entering any lines of code.

```
1. m = ConcreteModel()
2. m.fs = FlowsheetBlock(dynamic=False)
3. logging.basicConfig(filename='PyomoLog.log', encoding='utf-8', level=logging.INFO)
4.
5. #m.fs.propertiesR = GenericParameterBlock(**configuration)
```

```

6. fluidR = "propane"
7. m.fs.propertiesRC = HelmholtzParameterBlock(
8.   pure_component= fluidR,
9.   phase_presentation=PhaseType.MIX,
10.  state_vars=StateVars.PH,
11. )
12. m.fs.propertiesRC.default_enthalpy_mol_bounds = (-1000, 350000)
13.
14. m.fs.propertiesRW = HelmholtzParameterBlock(
15.   pure_component="h2o",
16.   phase_presentation=PhaseType.MIX,
17.   state_vars=StateVars.PH,
18. )
19.
20. #Compressor
21. m.fs.Comp = Compressor(
22.   property_package=m.fs.propertiesRC
23. )
24.
25. #Condenser
26. m.fs.Conds = HeatExchanger(
27.   delta_temperature_callback=delta_temperature_amtd_callback,
28.   hot_side_name="tube",
29.   cold_side_name="shell",
30.   tube={"property_package": m.fs.propertiesRC},
31.   shell={"property_package": m.fs.propertiesRW},)
32.
33. m.fs.Condc = HeatExchanger(
34.   delta_temperature_callback=delta_temperature_amtd_callback,
35.   hot_side_name="tube",
36.   cold_side_name="shell",
37.   tube={"property_package": m.fs.propertiesRC},
38.   shell={"property_package": m.fs.propertiesRW},)
39.

```

Figure 49 Excerpt from the heat pump model showing the steep learning curve required for IDAES

The Ahuora Platform simplifies the creation of the models by creating graphical representations for models like the code displayed in Figure 49. This takes a lot of the tedious model creating and constraint writing away from the user, thus simplifying model construction.

6.3.2 *Methods for extracting useful data for the PI&E method*

A central requirement of the PIE methodology is the ability to obtain accurate, structured, and relevant process data from the digital twin for downstream integration and targeting analysis. While the digital twin is built primarily for simulation and performance evaluation, its true value in process integration lies in its ability to serve as a reliable source of thermodynamically consistent data. This requires more than simply running the model it involves defining a clear data extraction strategy that captures the information necessary for heat recovery targeting, carbon footprint assessment, and techno-economic evaluation.

In this context, the simulation must be configured to automatically generate a stream data set containing all thermophysical properties and energy flows relevant to integration studies. This includes the identification and characterisation of hot and cold streams, their thermal capacities, temperature ranges, and associated duties. When the extracted stream data is to be used for pinch or exergy analysis, an additional requirement is that all thermal data be represented in a linear form suitable for composite curve construction. Pinch analysis assumes that the enthalpy change of each stream varies linearly with temperature over its specified range. In real process systems, however, this relationship is often non-linear, particularly in the presence of phase change or strongly temperature-dependent heat capacities. For example, evaporation and condensation occur at constant temperature, producing a vertical segment on the enthalpy–temperature diagram.

To address this, the digital twin must be configured to segment each stream into discrete, linear intervals, clearly separating sensible heating or cooling regions from latent heat transfer zones. While Aspen HYSYS provides automated data transfer capabilities to export stream data into Aspen Energy Analyzer, the built-in tool has limitations when applied to systems with significant phase change or specialised heating unit operations such as the DCH and DSI. In such cases, the automated segmentation may not accurately capture the true thermal requirements or operational characteristics of the process. Furthermore, Aspen Energy Analyzer does not support work targeting, a critical step in the PI&E methodology for evaluating and integrating mechanical work and power-driven technologies alongside thermal energy systems.

IDAES and DWSim does not provide a native, automated export function for pinch-compatible stream data.

As a result, an automatic export workflow was configured manually for each simulation platform. In the case of Aspen HYSYS and DWSim, this involved systematically reviewing the flowsheet and identifying streams to be exported, then (using the data linking tools) recording their temperature, heat duty, and enthalpy information into a linked worksheet. This worksheet was pre-formatted so that its contents could be directly copied into the Open Pinch Analysis workbook, where all pinch targets, composite curves, and grand composite curves could be generated. This manual selection and formatting ensured that only the relevant, linearised thermal data was extracted, while maintaining consistency with the process model and its thermodynamic definitions. Once this workflow was established, it became straightforward to run multiple cases and rapidly extract updated data which is an essential capability for supporting the iterative nature of the PI&E methodology.

For IDAES, the export process was significantly simpler. Again, key variables such as stream temperatures, heat duties, and enthalpies were identified within the flowsheet model and exported directly into a structured panda DataFrame. Because both IDAES and the Python implementation of Open Pinch are open-source and written in the same programming language, the two could be seamlessly linked. This allowed the simulation to be solved and pinch targets, composite curves, and grand composite curves to be generated automatically without the need for intermediate manual formatting. The result was a fully integrated workflow in which solving the process model directly triggered the generation of updated pinch analysis results, streamlining the iterative targeting and optimisation process within the PI&E methodology.

In addition to thermal data extraction, both simulation platforms were configured to automatically evaluate operating costs and carbon emissions as part of each model run. The methodology was consistent across HYSYS, DWSim and IDAES, with differences only in how the relevant process variables were accessed. In both cases, cost curve equations were linked to the model outputs, allowing electricity and fuel consumption to be translated directly into operating expenditure (OPEX) and associated CO₂ emissions. For HYSYS, the necessary variables such as electrical work for compressors and pumps, or thermal duty for fired heaters were gathered via the internal worksheets and referenced in the cost and emissions calculations. For IDAES, the same variables were accessed directly from the model using Pyomo expressions, eliminating the need for intermediate worksheets. This approach enabled fully automated evaluation: once a simulation was solved, the associated cost and carbon performance metrics were generated immediately. Integrating these calculations into the modelling workflow ensured that each scenario could be assessed not only for thermal integration

potential, but also for economic feasibility and environmental impact, both of which are essential to the PI&E methodology.

6.4 Results and Discussion

Utilising the methods and results established in Part A, all three simulation platforms (Aspen HYSYS, DWSIM, and the Ahuora platform powered by IDEAS) were successfully used to construct digital twins of a modern milk evaporator system. In each case, the model encompassed the full process boundary defined earlier, including preheating, heat treatment via direct steam injection (DSI) or direct contact heating (DCH), multiple-effect evaporation with vapour recompression, and downstream concentrate heating.

It is important to emphasise that, within the scope of this research, the definition of a digital twin extends beyond the creation of a steady-state process flowsheet. To qualify as a digital twin, a model must provide readily accessible data and support rapid iteration across scenarios and configurations. This ensures that the digital twin is not simply a static process model but rather a dynamic design and decision-support environment, capable of interfacing with optimisation workflows, electrification strategies, and integration analyses.

The following sections compare the three platforms across a set of technical and practical metrics, including model development, solver performance, thermodynamic fidelity, and integration potential. Together, these comparisons highlight both the strengths and limitations of each platform and illustrate the trade-offs between commercial ease-of-use, open-source accessibility, and advanced equation-oriented flexibility.

6.4.1 Comparison of digital twin results

The first essential question is whether the different platforms generate comparable results when applied to the same process. A digital twin must not only be flexible and transparent but also reliable in its predictions of mass balances, energy duties, and thermodynamic properties. In this section, the outputs from Aspen HYSYS, DWSIM, and the Ahuora platform powered by IDAES are compared to evaluate the degree of alignment across key performance indicators. This comparison provides confidence that methodological differences in modelling environments do not compromise the validity of the results, while also highlighting areas where deviations arise due to property package formulations, solver strategies, or unit operation representations.

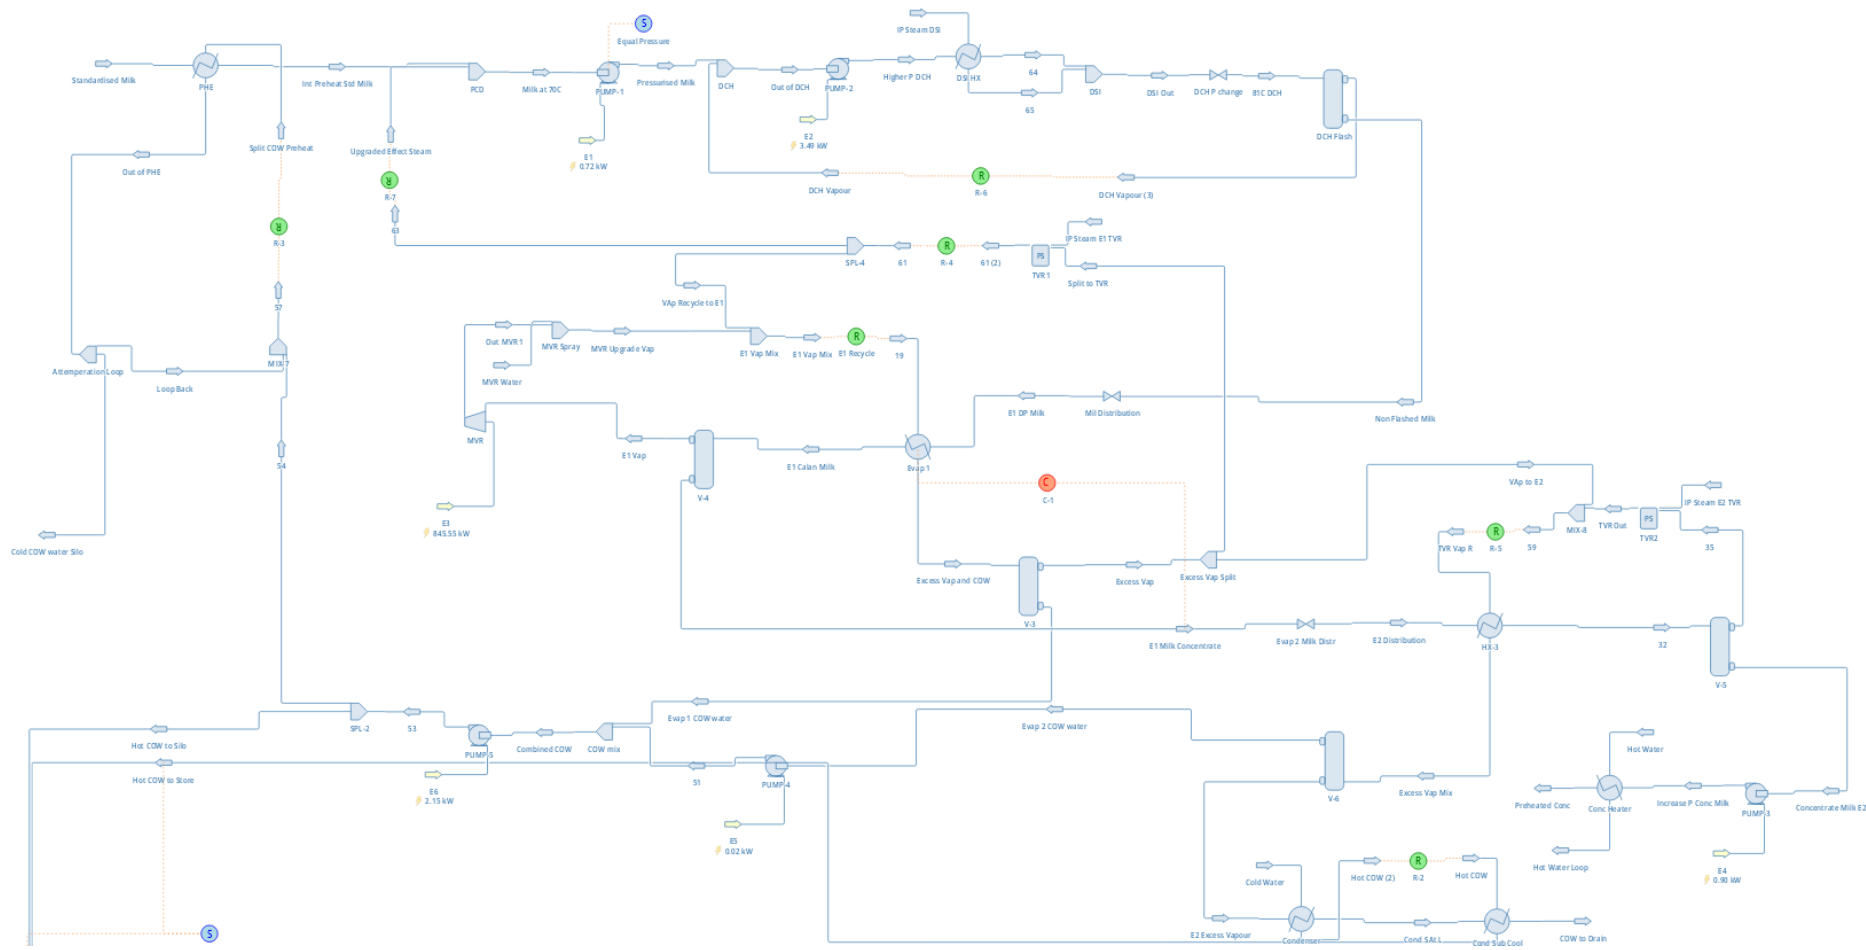


Figure 51 Flowsheet of one of the cases to create an electric milk evaporator in DWSim

To compare the alignment between different frameworks, comparison points were selected along key areas of the process as well as the time taken to solve from a standard non-converged state Table 18.

Table 18 Comparison of Values from different frameworks

Testing Location	Value from Hysys	Value from DWSim	Value from Ahuora
Amount of Heat Required across Preheat (Raw Milk to DSI)	22839 kW	23446 kW	22125 kW
MVR Work from Effect 1	1545 kW	1509 kW	1505 kW
Final Concentrate Heater (to 70°C)	259 kW	255 kW	242 kW
Time to Solve with Divergent Initial Values	6.47 s	73 s	12.43 s
Time to Solve with a Pre-converged Start	t > 1s	-	t > 1s

The results in Table 18 demonstrate that all three platforms were able to generate broadly consistent outcomes for the milk evaporation digital twin, though notable differences exist in both the predicted energy values and solver performance.

For the preheating stage (raw milk to DSI), all platforms estimated the heat duty within a relatively narrow range of 22.1–23.4 MW, a spread of around $\pm 3\%$. This variation primarily reflects differences in thermodynamic property package formulation. While the overall modelling approach was broadly similar, subtle factors introduced variability such as the smoothing functions required for equation-oriented VLE in Ahuora/IDAES, or the use of open-source food property data for fatty acids in DWSIM instead of hypothetical components.

In contrast, the predictions for MVR work from Effect 1 showed very close agreement across platforms. All three converged on values between 1505–1545 kW, a difference of less than 3%, indicating consistency where steam thermodynamics dominate.

For the final concentrate heater, differences were again small but slightly more pronounced. The Ahuora/IDAES platform predicted a duty of 242 kW, compared with 259 kW in HYSYS and 255 kW in DWSIM. Here, the higher concentration of milk solids in the stream likely accentuates differences in how each platform formulates and applies property models for concentrated, non-ideal fluids. These formulation differences contribute to the observed deviations, even though all three remain within an acceptable engineering range.

The most striking differences emerged in solver performance. HYSYS converged the flowsheet in 6.47 seconds from divergent initial conditions, while DWSIM required 73 seconds. The Ahuora/IDAES model achieved convergence in 12.43 seconds slower than HYSYS but an order of magnitude faster than DWSIM. This reflects the fundamental differences in solver architecture. HYSYS benefits from decades of commercial refinement and robust sequential-modular convergence routines, while DWSIM's open-source implementation shows slower recycle handling. The Ahuora/IDAES equation-oriented solver requires more mathematical setup but once structured, can converge efficiently and reliably. It is also worth noting that the creators of the ADTP are working on improving solving times.

6.4.2 Gathering data

A key requirement for the digital twin is that thermodynamic stream data can be extracted and transformed into a format suitable for process integration and exergy analysis. In this study, two distinct approaches were required depending on the simulation environment.

The screenshot shows the 'Spreadsheet' tab in HYSYS. At the top, there are tabs for 'Connections', 'Parameters', 'Formulas', 'Spreadsheet', 'Calculation Order', 'User Variables', and 'Notes'. Below these tabs is a 'Current Cell' section with a text box containing 'F5', a 'Variable:' label, an empty text box, an 'Exportable' checkbox, and an 'Angles in:' dropdown menu set to 'Rad'. There is also an 'Edit Rows/Columns' button. Below this is a spreadsheet table with 13 rows and 7 columns (A-F). The table contains the following data:

	A	B	C	D	E	F
1	Name	TS(C)	TT (C)	Mcp	H (kW)	Other Parameters
2	Cold Milk	8.000 C	95.00 C	2.500	2.283e+004 kW	2.470e+005 kg/h
3	Flash Vapour	86.98 C	86.48 C	0.5000	2061 kW	
4	2nd Flash Vapour	79.31 C	78.81 C	0.5000	1989 kW	
5	E1 Vapour Bleed	71.98 C	71.48 C	0.5000	3920 kW	
6	E2 Condenser	72.09 C	71.59 C	0.5000	374.3 kW	
7	Preheat COW 1	71.98 C	13.00 C	2.500	1.070e+004 kW	
8	Preheat COW 2	86.98 C	13.00 C	2.500	280.1 kW	
9	Preheat 3 COW	71.98 C	13.00 C	2.500	412.1 kW	
10	Preheat 4 COW	72.09 C	13.00 C	2.500	704.4 kW	
11	Preheat 5 COW	79.31 C	13.00 C	2.500	240.3 kW	
12	Vap COW Heat	72.09 C	13.00 C	2.500	38.91 kW	
13	Milk Concentrate	65.28 C	70.00 C	2.500	259.7 kW	

Figure 53 Automatic Spreadsheet Data Collection (Hysys)

For Aspen HYSYS and DWSIM, stream data were collected using linked workbooks within the simulation platform. These workbooks served as an intermediate layer between the simulator and the OpenPinch targeting tool. Hot and cold stream information including supply and target temperatures,

heat capacities, and duties were manually exported, structured into workbook templates, and then copied into the OpenPinch environment. While effective, this approach introduced additional workflow steps and a risk of manual error when reconciling large sets of process data. It also limited the speed of iteration, as any design change in the flowsheet required a repeat of the export–copy–paste cycle.

By contrast, the Ahuora platform powered by IDAES employed a direct data-frame connection to OpenPinch through Python. This integration bypassed the need for intermediate spreadsheets by streaming simulation results directly into structured pandas dataframes, which were then passed into the OpenPinch API. As a result, all streams were automatically formatted and immediately available for exergy and pinch analysis. This method enabled rapid iteration: updates to process conditions or unit operation configurations in IDAES were reflected instantly in the OpenPinch analysis without the need for manual intervention.

	Name	TS(C)	TT(C)	Mcp	H (kw)
0	Cold Milk	8.00	95.00	2.5	22830.00
1	Flash Vapour	86.98	86.48	0.5	2061.00
2	2nd Flash Vapour	79.31	78.81	0.5	1989.00
3	E1 Vapour Bleed	71.98	71.48	0.5	3920.00
4	E2 Condenser	72.09	71.59	0.5	374.30
5	Preheat COW 1	71.98	13.00	2.5	10700.00
6	Preheat COW 2	86.98	13.00	2.5	280.10
7	Preheat 3 COW	71.98	13.00	2.5	412.10
8	Preheat 4 COW	72.09	13.00	2.5	704.40
9	Preheat 5 COW	79.31	13.00	2.5	240.30
10	Vap COW Heat	72.09	13.00	2.5	38.91
11	Milk Concentrate	65.28	70.00	2.5	259.70

Figure 54 Automatic Spreadsheet Data Collection (Python IDAES)

While Aspen HYSYS includes an automated stream export tool, its application to this case study highlighted several limitations. As outlined in the methodology, the HYSYS automation is primarily designed for relatively simple petrochemical flowsheets, where unit operations are well defined and thermodynamic behaviour conforms to established property models. In the context of milk evaporation, this functionality encountered several challenges. First, much of the required thermodynamic data is embedded within black-box unit models, meaning that intermediate state information is not readily accessible. Second, the presence of composite streams, such as those involving vapour–liquid splits or direct steam injection, required manual reconciliation before export.

Third, exergy-relevant information, including stream enthalpy–temperature profiles or minimum approach temperatures, is not directly available from the automated tool.

A prototype for fully automated stream data collection is currently in beta development, to which the author has partially contributed too. This approach leverages Python-based dataframes and native OpenPinch integration, removing the need for manual intervention. By embedding stream queries within the simulation architecture itself, the prototype enables:

- Direct access to thermodynamic variables (temperature, enthalpy, entropy, flowrates) at every modelled state point.
- Customisable data formatting, ensuring outputs are directly compatible with integration tools.
- Real-time iteration, where design changes in the flowsheet propagate immediately into updated pinch or exergy analyses.

While this functionality is not yet fully production-ready, its development demonstrates the potential for equation-oriented, open-source frameworks to overcome long-standing bottlenecks in integration workflows. More importantly, it illustrates how the digital twin concept can move beyond “simulation plus export” towards a continuously connected environment, where modelling, targeting, and optimisation are all part of a single iterative loop.

Name	T _{supply} (°C)	T _{target} (°C)	ΔH (kW)	ΔT _{cont} (°C)	HTC (kW/m ² °C)
Chiller Water	8.00	4.00	456.00	0.50	1.00
CIP Water	15.00	84.00	630.00	2.50	1.00
Cow Heat	72.09	13.00	38.91	2.50	1.00
Direct Use	15.00	55.00	525.00	2.50	1.00
E1 vapour bleed	71.98	71.97	3920.49	0.50	1.00
E2 Condenser	72.00	71.90	364.05	0.50	1.00
HT Flash	86.98	86.88	2060.85	0.50	1.00
HT Flash 2	79.30	79.20	1988.77	0.50	1.00
Milk Concentrate	65.28	70.00	259.91	2.50	1.00
Preheat COW 1	71.98	13.00	10697.77	2.50	1.00
Preheat COW 2	88.88	13.00	279.79	2.50	1.00
Preheat COW 3	71.97	13.00	412.10	2.50	1.00
Preheat COW 4	71.90	13.00	704.42	2.50	1.00
Preheat COW 5	79.20	13.00	240.30	2.50	1.00
Raw Milk	8.00	95.00	22834.85	2.50	1.00

Figure 55 Prototype for Automatic Data Extraction combined with Targeting

From an industrial perspective, this shift is significant. Automating data collection not only reduces engineering effort but also minimises the risk of error, improves reproducibility, and enables more agile exploration of electrification strategies. The author's contribution to the Ahuora prototype reflects the broader research goal of embedding PI&E methods directly within digital twin environments, thereby lowering barriers for future industrial adoption.

6.4.3 Knowledge barriers

Whilst the creation of the digital twin was ultimately a technical success, delivering a model that is accurate, flexible, and fully capable of supporting the PI&E methodology. However, its development was far from trivial. Building such a system required a combination of deep chemical engineering expertise to ensure thermodynamic fidelity, realistic process representation, and operational relevance and advanced digital skills in coding, numerical optimisation, and data integration. These intersecting skill requirements present a barrier for many practising process engineers, particularly when working in open-source environments where user interfaces are less developed, documentation can be sparse, and modelling tasks demand a high degree of programming literacy. There is therefore both scope and an urgent need to lower these technological barriers, especially within open-source platforms, so that advanced digital twin strategies can be more readily adopted by the wider engineering community. Enhancing accessibility would not only broaden participation but also accelerate the application of PI&E methods in real industrial decarbonisation projects.

6.5 Conclusion

This chapter has demonstrated the creation and comparative evaluation of digital twins for a modern milk evaporator system using three simulation platforms: Aspen HYSYS, DWSIM, and the Ahuora platform powered by IDAES. Building on the design intent, boundary definition, and advanced property packages developed in earlier sections, the chapter established complete process models encompassing preheating, heat treatment, multiple-effect evaporation with vapour recompression, and final concentrate heating.

The results confirmed that all three platforms are capable of producing technically consistent digital twins, with key performance indicators such as heat duties and compressor work aligning within a narrow margin. Differences observed in preheating and concentrate heating stages reflected the sensitivity of results to thermodynamic property package formulation, particularly for concentrated milk streams where non-ideal behaviour is pronounced. Solver performance highlighted clear trade-

offs: HYSYS offered rapid convergence, DWSIM showed slower handling of recycles, and the Ahuora/IDAES framework solved quickly however could be challenging to initialise.

A central focus of this chapter was the extraction and structuring of thermodynamic stream data for use in process integration and electrification analysis. While HYSYS and DWSIM required intermediate workbooks and manual reconciliation, the Ahuora/IDAES environment enabled direct Python-based data exchange with OpenPinch, eliminating transcription steps and supporting real-time iteration. Furthermore, the development of a prototype for fully automated stream data collection, to which the author contributed, demonstrated the potential for continuously connected workflows in which modelling, integration, and optimisation are seamlessly linked. An avenue to be continued in the future.

The digital twin developed in this chapter thus not only establishes a robust representation of the milk evaporation system but also provides the foundation for the PI&E methodology applied in subsequent chapters. By enabling rapid, reliable, and transparent access to process data, the digital twin fulfils its role as both a simulation environment and a decision-support tool for electrification and decarbonisation of industrial process heat.

Part B: Process Integration and Electrification

In the second part of the thesis, attention shifts from the construction of the digital twin to its application as a platform for PI&E. Having established a validated and flexible representation of the milk evaporation system, the focus now turns to how this model can be used to systematically identify, evaluate, and design electrification pathways that minimise fossil fuel use and improve process efficiency.

Unlike Part A, which concentrated on the technical foundations of digitalisation and property modelling, Part B addresses the practical integration of electrification technologies within industrial processes. This includes operational optimisation, exergy- and energy-based targeting, and the incorporation of technologies such as mechanical vapour recompression (MVR), thermal vapour recompression (TVR), and industrial heat pumps. Both greenfield and retrofit contexts are considered, recognising that while greenfield designs provide the opportunity for system-wide reconfiguration, the majority of industrial decarbonisation must be achieved through retrofit pathways constrained by existing equipment, utilities, and site integration.

The chapters in this part demonstrate how the PI&E methodology can be applied iteratively, progressing from baseline optimisation to advanced targeting, and from conceptual design to techno-economic evaluation. Case studies are used to illustrate the methodology in practice, highlighting the trade-offs between thermodynamic performance, operability, and cost. In doing so, Part B positions the digital twin not as an isolated modelling exercise, but as a critical decision-support tool for guiding real-world industrial decarbonisation strategies.

Chapter 7 Process Integration and Electrification for Greenfield Design

7.1 Introduction

Building on the foundation of the design DT, the present chapter shifts the focus from enabling capability to applied design methodology. Specifically, it demonstrates how the PI&E framework can be deployed in a greenfield context, where process layouts are unconstrained by legacy infrastructure and electrification technologies can be embedded from the outset. Greenfield design offers unique opportunities to optimise the placement and integration of high-efficiency, electrically driven technologies such as mechanical vapour recompression (MVR), high-temperature heat pumps, and potentially other technologies within fully integrated heat recovery networks.

While conventional Process Integration techniques aim to minimise thermal utility consumption, they rarely account for the thermodynamic and operational characteristics that govern the performance of electrification technologies, such as temperature lift penalties, part-load behaviour, or integration with process control requirements. The PI&E method used here extends traditional approaches by incorporating exergy-aware targeting, iterative process simulation, and heat exchanger network (HEN) synthesis, ensuring that electricity use is minimised while operational feasibility is preserved.

This chapter applies the method to the design of a modern milk evaporation system which is a representative high-energy process in New Zealand's dairy industry. Starting from the validated property packages and digital twin environment, the methodology progresses through base-case simulation, electrification targeting, technology placement, and HEN optimisation, culminating in a fully electrified, energy-efficient process design. The resulting case study not only validates the PI&E approach for greenfield applications but also offers a replicable blueprint for other low- to medium-temperature industries pursuing deep decarbonisation through electrification.

The decarbonisation of industrial process heat is both a critical and complex challenge for achieving national and global climate targets. In greenfield plant design, where process configurations can be defined from the outset, there is a unique opportunity to embed low-emissions technologies directly into the process architecture rather than relying on later retrofits. This proactive approach allows engineers to integrate electrification technologies such as high-temperature heat pumps, mechanical vapour recompression (MVR), and hybrid thermal–electric utility systems within optimised thermal

networks from the very beginning, maximising both thermodynamic efficiency and economic performance.

In sectors such as dairy processing, where evaporation and drying dominate site-wide energy demand, process electrification offers the potential for substantial reductions in greenhouse gas emissions while maintaining product quality and throughput. The renewable-rich electricity mix in New Zealand makes this opportunity particularly compelling, enabling the displacement of fossil-fuelled steam generation with clean electric utilities. However, capitalising on this potential requires a design methodology that extends beyond conventional heat recovery targeting. Traditional Process Integration (PI) techniques focus on minimising thermal utility demand, but do not explicitly account for the work–heat trade-offs, performance characteristics, and operational constraints of electrically driven technologies.

The significant contributions of this chapter, within the context of the overall thesis, are:

- Synthesis of a new PI&E method specifically adapted for greenfield applications, enabling the iterative development of fully electrified process systems with optimal placement of heat recovery and electrification technologies.
- Development of an original, fully electric milk evaporation system design that demonstrates lower operating costs and improved efficiency compared to flowsheets reported in the literature, while satisfying industrial operability requirements.
- A sensitivity analysis of the evaporation system simulation to evaluate the robustness of the design across a range of process operating parameters, providing insights into scalability and adaptability for other industrial contexts.

These contributions collectively provide both a methodological advancement in PI&E and a practical case study that can serve as a replicable blueprint for other low- to medium-temperature process industries seeking deep decarbonisation through electrification.

7.2 Methods

The greenfield application of the (PI&E) methodology follows the staged framework established in Chapter 3. In this context, the absence of legacy constraints allows the targeting, simulation, and

integration steps to be pursued iteratively from first principles, with each iteration informed by both thermodynamic limits and practical operability considerations.

These constraints include safety, product quality, operating cost, and process control and operability requirements. Minimising electricity use is achieved by maximising the heat recovery within the process and the appropriate placement of vapour compression with the minimum temperature lift.

The overall process is illustrated in Figure 56 and comprises the following core stages:

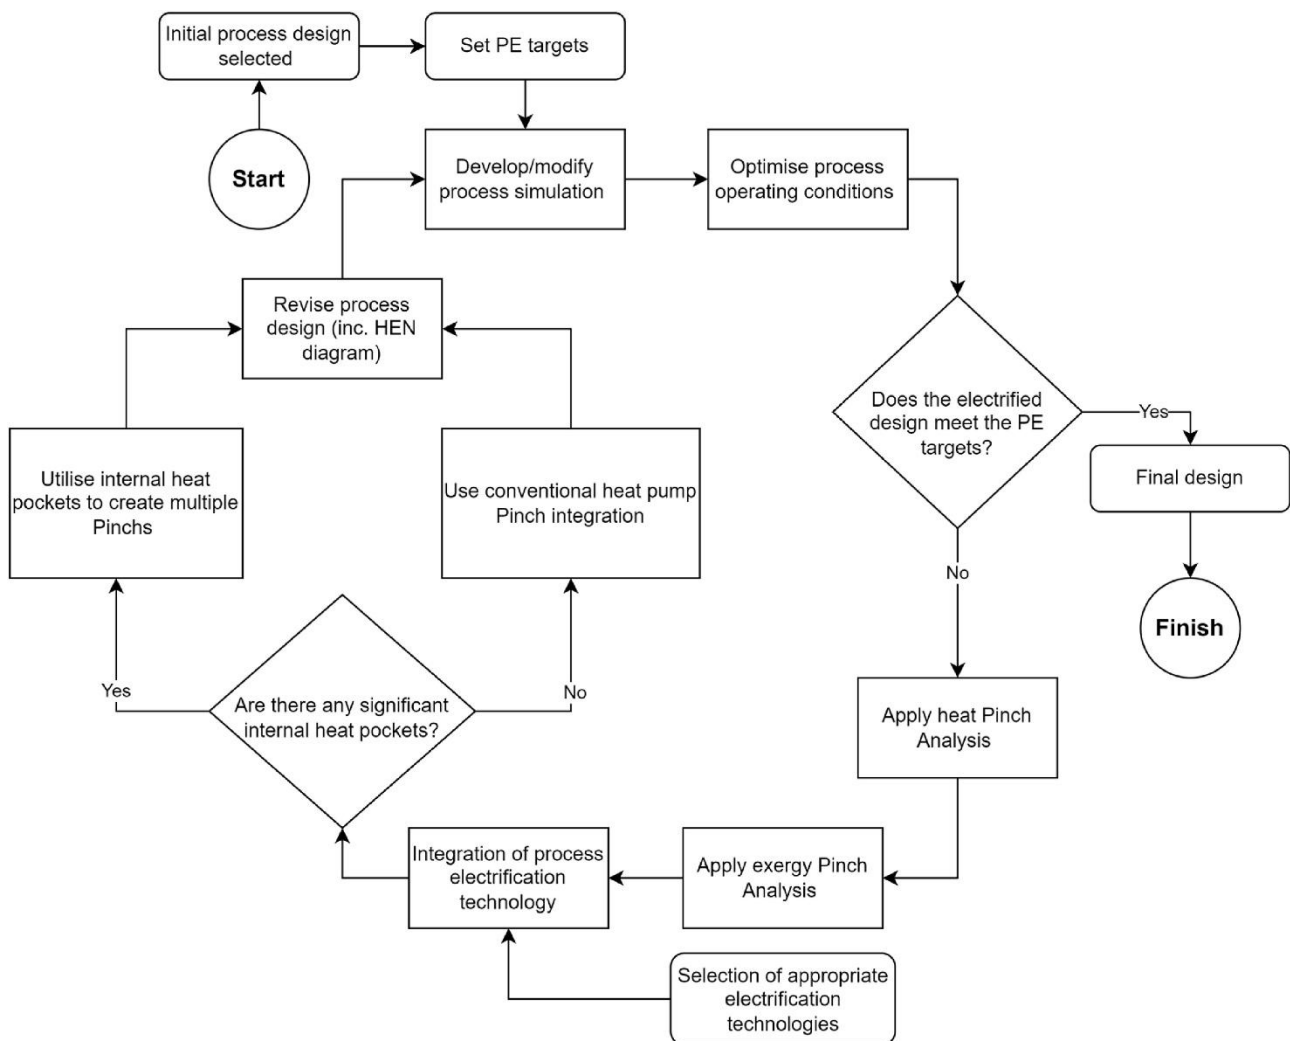


Figure 56 Process Integration and Electrification (PI&E) method.

7.2.1 Base-Case definition, Electrification Target and Digital Twin Development

The method begins with the construction and simulation of a base-case design for the target system, which in this study is a modern milk evaporation process representative of large-scale New Zealand dairy operations. This base-case serves as the reference configuration from which all subsequent modifications are developed, evaluated, and compared. Each of the subsequent methodological steps is applied sequentially and iteratively, with the process continuing until a desired level of electrification is achieved. This level can be described as an electrification substitution target (i.e. 0-100%) and/or a minimum COP required for the electric technology to be a viable design.

Although the overarching stages of the method are presented in a logical order, as illustrated in Figure 56, valuable design insights can emerge at any stage, sometimes prompting deviations from the nominal sequence to refine or re-evaluate earlier steps. To ensure that the investigated configurations are not only thermodynamically optimal but also operationally and commercially viable, the design process was conducted in continuous dialogue with industry stakeholders. Input from local milk processing companies and external experts was used to verify the practicality, operability, and acceptability of proposed designs within the real-world context of dairy processing.

7.2.2 Design Digital Twin and evaluation

See Chapter 4, 5 and 6 for a detailed description of creating and evaluating a digital twin

7.2.3 Targeting (Pinch Analysis)

The pinch analysis (PA) for the base-case milk evaporation system was conducted using the method described by Walmsley (2016), which is itself an adaptation of the approach developed by Westphalen & Maciel (2000). This method is particularly suited to food and dairy applications, where both process and phase-change heating and cooling loads must be accurately represented. In conventional evaporator studies, it is common to present the grand composite curve (GCC) for background processes separately from the evaporation and condensation heat loads of each effect Walmsley (2016). While this separation can clarify the thermodynamic structure of the system, the initial stage of greenfield process design benefits from a more holistic approach. In particular, including the evaporation loads within the site-wide GCC provides a complete picture of the thermal

demand profile and allows for the early identification of integration opportunities for mechanical vapour recompression (MVR) and other electrification technologies.

Pinch Analysis (PA) is a thermodynamic targeting method used to determine the minimum heating and cooling requirements for a given process configuration based solely on-stream temperature and heat load data. Hot and cold streams are shifted by half of the specified minimum temperature difference (ΔT_{\min}) to create composite curves representing the net heat availability and demand across the temperature spectrum. The point at which the shifted curves approach most closely defines the pinch temperature, which represents the boundary between heat recovery above and below the pinch. From these curves, the minimum hot and cold utility targets can be calculated without designing the heat exchanger network.

To implement PA the user is required to define the supply and target temperatures, heat duties, and individual minimum temperature difference contributions (ΔT_{cont}) for each stream. For the evaporation system, ΔT_{cont} values were assigned as 2.5 °C for liquid heat exchange and 0.5 °C for vapour/condensate exchange, reflecting industrial practice in the dairy sector. These are then added to the Problem Table Algorithm (PTA). Hot streams (i) are shifted down in temperature, while cold streams (j) are shifted up in temperature.

$$T_{\text{hot},i}^* = T_{\text{hot},i} - \Delta T_{\text{cont},i} \quad \text{Equation 49}$$

$$T_{\text{cold},j}^* = T_{\text{cold},j} + \Delta T_{\text{cont},j} \quad \text{Equation 50}$$

From the PTA, the minimum hot and cold utility requirements and the associated pinch temperatures were determined. The resulting grand composite curve (GCC) was then used to identify thermodynamically feasible points for integrating high-efficiency heat pumps, MVR systems, and other electrification technologies.

7.2.4 Targeting (Exergy Pinch Analysis)

While pinch analysis offers a robust framework for establishing the thermodynamic limits of heat recovery, it does not fully account for the interdependence between thermal integration and the deployment of electrification technologies such as heat pumps and mechanical vapour recompression (MVR). These technologies convert mechanical or electrical work into useful thermal energy, with

performance characteristics that are strongly dependent on temperature lift, heat source and sink quality, and the broader process integration strategy.

In the context of industrial decarbonisation, electrification has become an increasingly critical design consideration, necessitating analytical tools capable of evaluating both the quantity and quality of energy flows. Exergy-based targeting provides such capability by quantifying the maximum useful work obtainable from each energy stream and identifying the irreversibility's associated with thermal and mechanical transformations. Subsequently new Exergy Pinch Analysis (EPA) tools and methods are proposed that can set lower bound work targets by acutely balancing process heat recovery and heat pumping.

The exergy evaluation and targeting procedures employed in this study follow the principles of the ExPAND method (Aspelund et al., 2007) and the graphical representation introduced by Marmolejo-Correa & Gundersen (2013). Furthermore, the extension developed by Hamsani et al. (2018) demonstrates how an exergy-based grand composite curve (GCC) can be used to determine accurate targets for the total exergy deficit and surplus within a process system. Readers seeking further theoretical context are encouraged to consult these two latter works.

It is also important to note that Exergy Targeting been a core tenant of the PI&E method however further expansions and communications have been made about this topic in other publications predominantly Walmsley et al. (2024) to which the author contributed to. The method has been implemented in the OpenPinch Open-Sourced Tool.

7.2.4.1 Stream Data Extraction

Accurate and consistent stream data is essential exergy-based targeting as well. In this study, hot and cold process streams were extracted directly from the validated base-case digital twin to ensure complete alignment between the simulation model and the integration analysis. This approach eliminates the inconsistencies that can arise when stream data is compiled from disparate sources or simplified process representations.

Table 19 Representation of Gathering Stream Data

Steam Name	T supply (°C)	T Target (°C)	ΔH (kW)
Cold Stream One	70	90	100
Hot Stream One	70	69	300

7.2.4.2 Initial Stream Temperature Shifts and the Problem Table

Following the extraction of process stream data, the next step is again similar to PA which is to apply temperature shifts to account for the ΔT_{cont} associated with each stream. This step ensures that the thermal driving forces required for realistic heat exchanger operation are embedded in the targeting analysis. The temperature shifts are applied to create an intermediate temperature scale (T^*), which and like PA enables the PTA to be applied consistently across all streams.

Once the temperature shifts have been applied, the PTA is used to evaluate the process across defined temperature (or exergetic temperature) intervals.




Figure 57 Representation of Shifting Stream Temperatures

7.2.4.3 Heat Recovery Pocket Analysis and Cutting

This is where the EPA start to deviate from the conventional PA. The grand composite curve (GCC) provides a visual representation of process heat surpluses and deficits across the temperature spectrum, enabling the identification of “heat recovery pockets” (this is defined as regions where internal process heat recovery is thermodynamically possible). In conventional process integration, particularly in Total Site Heat Integration, these pockets are typically removed from the GCC before

further analysis. This approach simplifies the subsequent targeting process and reflects the assumption that pockets are best utilised for direct process-to-process heat recovery. However, when large temperature differences exist within a pocket, this convention can overlook opportunities to harness the thermodynamic potential of these surpluses through heat engines, heat pumps, or hybrid configurations, thereby improving overall process efficiency.

Subsequently previous work by Hamsani et al. (2018) assumed that all pockets were dedicated solely to internal heat recovery, an approach that simplifies the exergy analysis but also removes the possibility of exploiting pockets for additional thermodynamic efficiency gains. To address this additional pocket-cutting strategies are considered, each representing a different internal heat exchanger network configuration:

1. Full Cut: The entire pocket is removed from further analysis, as in conventional Total Site Heat Integration. This eliminates the pocket's thermal and exergetic content from the targeting stage, effectively assuming it is fully used for direct process heat recovery.
2. Minimum Cut (Min-Cut): The pocket is retained in the analysis to the greatest extent possible, subject to the minimum approach temperature between process and utility ($\Delta T_{min,U}$). This approach provides the theoretical upper bound for heat recovery potential and establishes a thermodynamic benchmark for exergy and work targets.

$$\Delta T_{min,cut} = \Delta T_{min,U} \quad \text{Equation 51}$$

3. Partial Cut: Only the surplus and deficit segments of the pocket within a user defined $\Delta T_{min,cut}$ are removed. This intermediate strategy allows for variable degrees of pocket utilisation and enables investigation of how different $\Delta T_{min,cut}$ values influence work and exergy targets.

It must be noted that the minimum cut must always be greater than the minimum approach temperature between process and utility $\Delta T_{min,cut} > \Delta T_{min,U}$. The Full-Cut approach sets ΔT_{cut} to the maximum ΔT of the largest pocket.

For the purposes of this thesis only the full pocket cut strategy will be implemented. The other methods for the electric milk evaporator can be found in Walmsley et al. (2024)




Figure 58 Representation of Pocket Cutting Strategies

7.2.4.4 Calculate the Energetic temperature

Previous studies, including Hamsani et al. (2018) have employed an exergy-based GCC to represent process integration potential. While this approach is effective when all process streams operate entirely above or entirely below the dead state temperature, difficulties arise when streams cross the dead state. Above ambient conditions, hot streams act as exergy sources and cold streams as exergy sinks; however, below ambient, these roles are reversed. Cold streams become exergy sources and hot streams become exergy sinks (e.g., in refrigeration systems). When streams cross this boundary, a single exergy GCC cannot clearly display both regions without becoming difficult to interpret.




Figure 59 Representative Net Heat Load Curve

To address this, the present study adopts the concept of net exergy load curves (NXLCs), which explicitly separate sources and sinks into two distinct profiles. To construct NXLCs, each heat load segment from the modified GCC is first converted into its corresponding exergetic temperature and exergy flow. Using the shifted supply (s) and target (t) temperatures (T_s , T_t), the corresponding exergetic temperatures (θ_s , θ_t) are calculated as:

$$\theta = T_0 \left[\frac{T^*}{T_0} - \ln \left(\frac{T^*}{T_0} \right) - 1 \right] \quad \text{Equation 52}$$

The temperature-based exergy change (ΔX) for each segment is then determined as:

$$\Delta X = C_p(\theta_s - \theta_t) \quad \text{Equation 53}$$

where T_0 is the dead state temperature in Kelvin, and C_p is the heat capacity flow rate ($\dot{m} c_p$). Equation 53 evaluates only the temperature-dependent component of exergy, which in the case of this study is a reasonable approximation of total thermomechanical exergy because pressure effects are negligible.

With these calculated values, the PTA is applied separately to all heat surplus segments (sources) and to all heat deficit segments (sinks), producing an exergy-based problem table and the corresponding NXLCs. The dead state reference conditions suggested are 15 °C (288.15 K) and 101.325 kPa, representative of typical ambient conditions.

The NXLC approach presented in this work is an original contribution of this study. By explicitly separating exergy sources and sinks, it enables clearer identification of work recovery and work input opportunities in processes that operate across the dead state temperature. These conditions are common in electrified industrial systems.

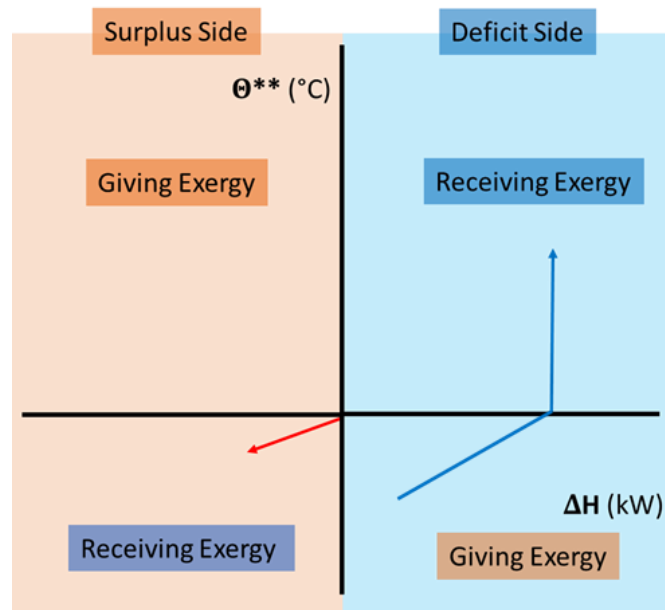


Figure 60 Representative Net Exergy Load Curve

For processes operating across the dead state temperature (i.e., ambient temperature), the designation of exergy sources and sinks is reversed depending on the temperature region. Above ambient temperature, hot streams function as exergy sources and cold streams as exergy sinks. Conversely, below ambient temperature, hot streams act as exergy sinks (as in refrigeration systems), while cold streams serve as exergy sources. In this study, the x-axis of the NXLC plot is used to distinguish between these two regions, with the origin corresponding to the dead state temperature. Exegetic temperatures are expressed as positive values regardless of whether they occur above or below ambient, ensuring consistency in the calculation of exergy differences.

7.2.5 Net shaft-work calculation

The final stage of the methodology is to determine the net shaft-work target for the process, based on the exergy surplus and deficits the principles of reversible heat engines and heat pumps. In a reversible (Carnot) heat engine operating between temperature levels T_H and T_L , entropy generation is zero and there is no exergy destruction. Under these ideal conditions, the work output is given by:

$$W_{HE} = W_H - W_L = X_H - X_L \quad \text{Equation 54}$$

where W_{HE} is the reversible work generated, W_H is the work transferred from T_H , W_L is the work transferred from T_L and X is the corresponding exergy flow. For practical systems, the work is reduced by the second law (Carnot) efficiency (η_{II}) giving:

$$W_{HE,act} = \eta_{II,HE}(X_H - X_L) \quad \text{Equation 55}$$

A similar derivation applies to heat pumps, where the actual work input

$$W_{HP,act} = \frac{1}{\eta_{II,HP}}(X_H - X_L) \quad \text{Equation 56}$$

Hamsani et al. (2018) combined these relationships to generate a net shaft-work target for a process:

$$W_{net(high)} = W_{comp} - W_{turb} \approx \frac{1}{\eta}X_{rej} - \eta X_{req} \quad \text{Equation 57}$$

where W_{comp} is the work of compression, W_{turb} is the work recovered through expansion, X_{rej} is the exergy rejection (below the pinch), X_{req} is the exergy requirement (above the pinch), and $\eta = \eta_{II,HP} = \eta_{II,HE}$. A second-law efficiency of 50% is typical for both practical heat engines and heat pumps (Arpagaus et al., 2018), making this a reasonable design assumption, although other values can be used to explore sensitivity.

In this method, the pockets on the GCC are cut according to the defined strategy before the exergy analysis. Equation 57 assumes that all exergy surpluses are used in heat engines and all exergy deficits are supplied by heat pumps. This can overestimate the actual work requirement, as a more efficient approach is to use exergy surpluses directly to drive heat pumps wherever possible, avoiding the compounded inefficiencies of converting surplus heat to work via a heat engine and then using that work to drive a heat pump.

To reflect this improved strategy, the new method introduces the parameter γ , defined as the fraction of exergy surplus segments (S) utilised directly by heat pumps to meet exergy deficits (D). Any remaining surplus is recovered by a heat engine. As a result, this study defines γ as the fraction of heat surplus segments supplying heat pump devices (as opposed to heat engine devices) The net shaft-work target is therefore calculated as:

$$W_{\text{net}} = \frac{1}{\eta_{\text{II,HP}}} (\sum \Delta X_D - \gamma \sum \Delta X_S) - \eta_{\text{II,HE}} (1 - \gamma) \sum \Delta X_S \quad \text{Equation 58}$$

subject to:

$$0 \leq \gamma \leq \min \left(\frac{\sum \Delta X_D}{\sum \Delta X_S}, 1 \right) \quad \text{Equation 59}$$

where ΔX_D is an exergy deficit segment on the NXLC, and ΔX_S is an exergy surplus segment.

In the Full-Cut pocket-cutting strategy, all exergy deficits are above the pinch, and all exergy surpluses are below the pinch, making $\gamma = 0$ and recovering the formulation of Equation 57. At the other extreme, setting $\gamma = \min \left(\frac{\sum \Delta X_D}{\sum \Delta X_S}, 1 \right)$ yields a lower-bound net shaft-work target for an ideally designed system in which heat pumps and heat engines are perfectly matched to the available surpluses and deficits, given the chosen minimum approach temperatures, pocket-cutting strategy, and assumed Carnot efficiencies.

Again, for the purposes of the thesis only the Full-Cut strategy was considered. Work targets for the other strategies can be found in Walmsley et al. (2024).

7.2.6 Operating set-point optimisation

Once energy and exergy targets have been established, operating set-point optimisation can be performed to further close the gap between the current process performance and the thermodynamic limits (these inefficiencies can be identified during targeting). Many process variables such as temperatures, pressures, and recirculation rates are “soft” parameters that provide degrees of freedom for constrained optimisation aimed at minimising total work (electricity consumption) while maintaining product quality.

Adjusting these parameters can alter the thermal profile of the process, thereby changing both the quantity and quality of recoverable heat. Such modifications often have interdependent effects across the flowsheet, influencing utility demand, heat pump temperature lifts, and pinch locations. For this reason, optimisation must be carried out within the digital twin environment to ensure that all mass and energy interactions are accurately captured.

Following optimisation, a revised set of process stream data is extracted and reanalysed using pinch analysis (and, where relevant, exergy-based targeting) to update heat integration targets and pinch temperatures. This iterative feedback between targeting, simulation, and operational adjustments ensures that subsequent technology selection and integration strategies are based on the most efficient achievable operating conditions.

7.2.7 Process electrification technology selection and integration

Process electrification can be achieved through several pathways, each involving the replacement of conventional fossil-fuel-based thermal utilities with electrically driven technologies. These include:

- Processing equipment redesign to directly utilise electricity in place of steam, such as ohmic heating systems.
- Closed-cycle heat pumps, typically based on standard vapour-compression technology.
- Open-cycle heat pumps, such as mechanical vapour recompression (MVR) systems, which directly upgrade process vapour streams.
- Electrified processing technologies, for example, pulsed electric field (PEF) systems (Jeyamkondan et al., 1999), which introduce entirely new process mechanisms enabled by electrical energy.

Standard heat pump technologies are generally capable of supplying heat at temperatures up to approximately 100°C, while commercial high-temperature and very-high-temperature heat pumps can deliver supply temperatures of up to 165°C (Schlosser et al., 2020a). The COP is a critical performance indicator for any heat pump system and is strongly influenced by the mean temperature lift (ΔT_{lift}) between the heat source and the heat sink. As ΔT_{lift} increases, the COP typically decreases, reducing overall system efficiency. In addition to achieving the appropriate temperature lift, correctly matching the thermal profiles of heat sources and sinks is essential to maximise both energy recovery and operational stability.

Accurate COP estimation is essential for both the thermodynamic integration of heat pumps into process heat recovery networks and for the techno-economic assessment of their feasibility. Over- or underestimation of COP can lead to suboptimal design decisions, either overstating potential energy savings or prematurely dismissing viable integration opportunities. Consequently, technology

selection in this study is informed by detailed performance modelling of candidate systems, aligned with the targeted temperature levels, available heat sources, and required process heat duties identified in earlier sections.

7.2.8 Conventional and “Cross-Pocket” Heat Pump integration principles

The appropriate placement principle for heat pumps defines the optimal positioning of the heat source and sink to achieve maximum utility reduction. In straightforward process configurations, applying this principle typically results in the highest achievable efficiency. However, in systems with significant opportunities for internal heat recovery, which are represented as large “heat pockets” on the grand composite curve (GCC) additional integration strategies become possible. In such cases, process electrification (PE) design must balance a trade-off between reducing the mean temperature lift (ΔT_{lift}), which improves operational efficiency, and minimising the number of heat pump units, which influences capital cost.

Figure 61 illustrates two contrasting approaches to heat pump placement for full process electrification. In both rows, the first image shows the placement of the initial heat pump. The second image presents the GCC after removal of the heat load segments associated with this first unit, with the placement of a second heat pump. The third image displays the original GCC with both heat pumps integrated.

The top row (a-c) demonstrates the conventional approach, in which the heat pocket is “cut” and the cross-pinch (or appropriate placement) principle is applied to fully electrify the heat supply. This method prioritises adherence to the established placement principle and typically minimises the number of units required.

The bottom row presents a novel *cross-pocket integration* approach. Here, the heat pocket itself is exploited as part of the integration strategy, potentially enabling higher heat pump coefficients of performance (COPs) by reducing the mean ΔT_{lift} . While the placement of the first heat pump in this method appears to contravene the conventional appropriate placement principle, the approach can yield efficiency gains in situations where lowering the temperature lift outweighs the impact of introducing additional units. For full electrification, minimising the mean ΔT_{lift} remains the primary driver of COP, but cross-pocket integration may necessitate a greater number of heat pump units to meet the deficits created within the pocket.

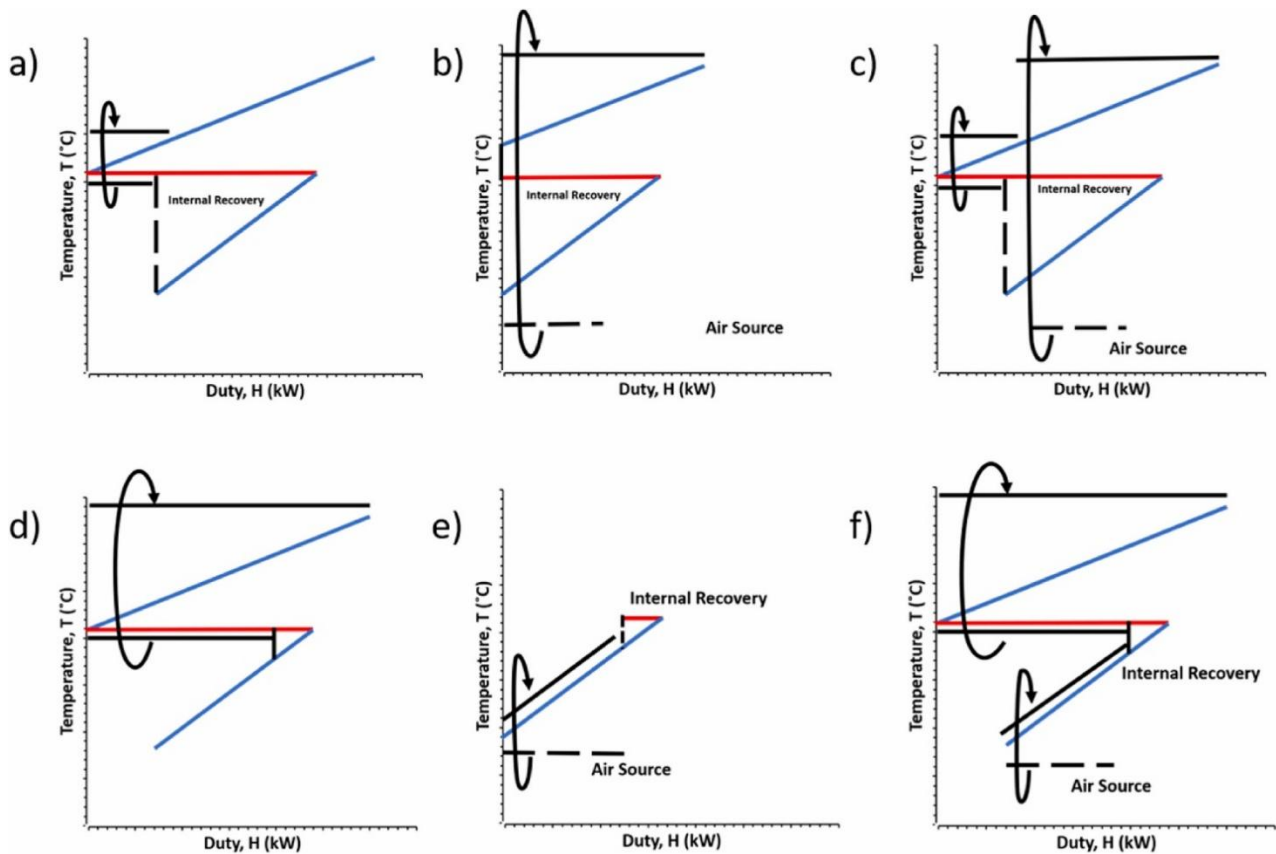


Figure 61 (a–c) Traditional air source heat pump with low COP overall, (d–f) Utilisation of a heat pocket to produce a higher COP.

In summary, the choice between conventional and cross-pocket integration strategies depends on the specific thermal profile of the process, the size and location of heat pockets, and the relative weighting of operational efficiency versus capital expenditure in the design objectives. This is where the evaluation step is critical.

7.2.9 Heat Exchanger Network design

Once the heat pump(s) have been placed and the corresponding source and sink stream segments identified, the heat exchanger network (HEN) must be designed or modified to accommodate the new configuration. The grid diagram and Pinch Design Method (PDM) developed by Linnhoff & Hindmarsh (1983) provide a systematic framework for designing the HEN for the remaining heat recovery problem. This approach ensures that the network satisfies both the minimum utility requirements established during targeting and the integration constraints associated with the electrification technologies.

With the HEN configuration defined, the complete process system (including all electrification units, utility systems, and heat recovery loops) can be modelled back into the digital twin environment.

7.2.10 Iterative design approach

To reiterate the defining feature of the Process Integration and Electrification (PI&E) methodology is its inherently iterative nature. Each methodological step from targeting and operating set-point optimisation, through technology selection and integration, to heat exchanger network (HEN) design, can alter the thermal, exergetic, or economic profile of the process. These changes often create new opportunities for improvement or reveal constraints that necessitate revisiting earlier decisions.

Iteration ensures that the design progressively converges toward a configuration that balances thermodynamic efficiency, operational feasibility, and economic viability. For example, modifications in HEN design may affect heat pump temperature lifts, prompting a re-evaluation of technology selection; similarly, optimised operating conditions may alter pinch locations, requiring updated heat recovery targets. By maintaining a feedback loop between analysis, simulation, and design, the PI&E method avoids suboptimal “one-pass” solutions and instead evolves toward the best practicable electrified process configuration.

An additional advantage of the iterative approach is that it can form the basis for a staged implementation plan. If the complete electrification solution is prohibitively expensive to implement in a single step, an earlier iteration of the design may still deliver substantial energy and emissions reductions while requiring lower capital investment. This flexibility allows decision-makers to adopt electrification in manageable stages, balancing short-term constraints with long-term decarbonisation goals.

7.3 Milk Evaporation Case Study Refresh

As described in Chapter three (and subsequent chapters) the case study used to display the method is a milk evaporation system. The base case design (Figure 62) is a conventional two-effect evaporation system, with a modern MVR and TVR structure and operation based on the work of Walmsley et al (2016). Please refer to Chapter four a detailed explanation of the case study.

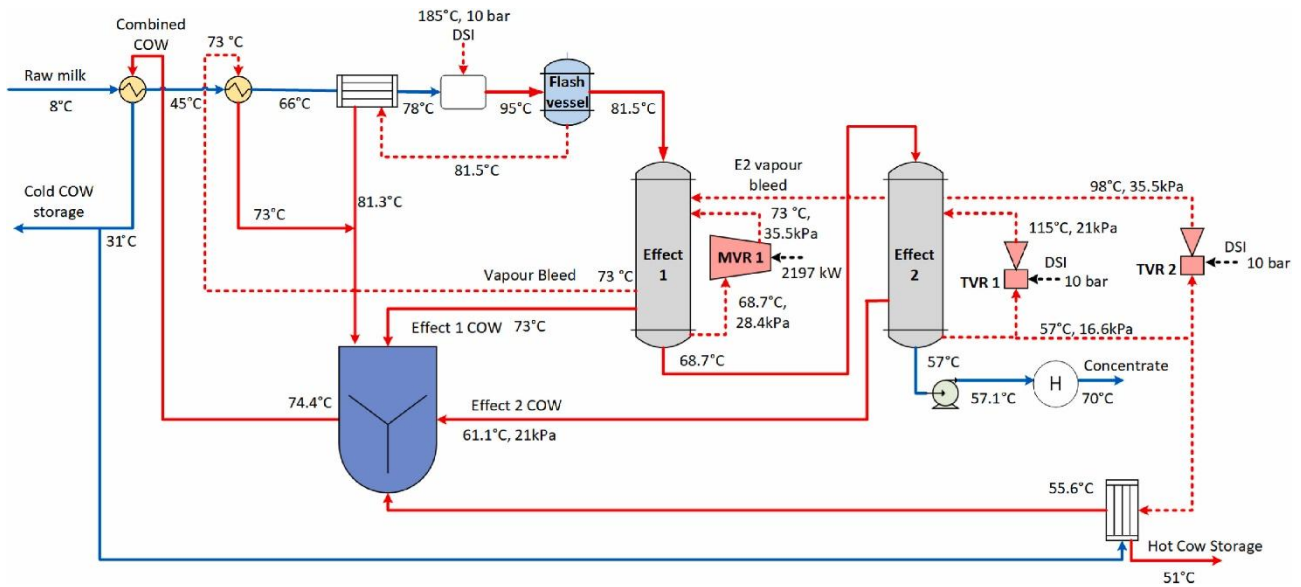


Figure 62 The base case design of a two-effect evaporator with MVR and TVR on each effect.

In the base-case configuration, raw milk is first preheated to 45° using a combination of condensate streams, collectively referred to as combined COW (Condensate of Whey) water. The vapour generated from the first evaporator effect (vapour bleed) is then utilised to further increase the milk temperature to 66°C. To preserve product quality, it is desirable to minimise the residence time and contact time of the milk with process equipment in this temperature range. Accordingly, direct steam injection (DSI) is employed to rapidly raise the milk temperature to 95°C, after which flash cooling reduces the temperature to 85°C while generating low-pressure vapour suitable for heat recovery. In the New Zealand dairy industry, DSI is often specified due to its beneficial effect on milk flavour.

The milk feed rate through the evaporation system is scaled to produce sufficient concentrate for the manufacture of 30 t/h of milk powder product. For comparative analysis, Specific Energy Consumption values are expressed in terms of both powder production (kWh/t_p) (which are calculated by dividing the total heat supplied by the quantity of powder exiting the spray dryer) and water removed (kWh/t_w), providing a complementary measure of thermal efficiency.

Mechanical vapour recompression (MVR) integrated evaporators derive limited benefit from increasing the number of effects, as they do not rely on the vapour cascading principle fundamental to traditional multi-effect evaporators. Nevertheless, for milk evaporation, inclusion of a second effect has become standard practice. As the milk becomes more concentrated, its boiling point elevation increases, necessitating a higher temperature driving force to sustain evaporation rates. A second effect also provides additional operational flexibility to control the final solids concentration in the

product stream. For these reasons, the present study focuses on a two-effect MVR milk evaporator configuration. While future work could examine systems with a greater number of effects, any potential efficiency improvements are expected to be marginal relative to the additional capital investment required.

7.3.1 Solution constraints and ideals

The development of the electric milk evaporator design itself was constrained by industry-standard processing requirements. Specifically, the requirements included:

- No direct heat exchange between refrigerants and milk streams to avoid product contamination with even the smallest of leaks,
- High heating rates in the milk heat treatment section (milk at $>80\text{ }^{\circ}\text{C}$) ideally using direct contact heating, and
- Heating the milk through the $45\text{ }^{\circ}\text{C}$ – $60\text{ }^{\circ}\text{C}$ range in one heat exchanger to prevent thermophile growth.

Ideals for the design included:

- Minimising the amount of electricity consumption while keeping within practical constraints,
- Simplifying the process as much as possible to minimise capital and footprint, and
- Where possible, localising the heat integration and heat pumps for better operability and control.

In addition, variations in operating parameters, such as shell-side and tube-side pressures in the effects, is possible to minimise the overall work required of the system.

7.4 Results

The results follow the proposed PI&E method to demonstrate its practical application to a milk evaporation system and its associated clean-in-place (CIP) system.

7.4.1 Initial Pinch Analysis, simulation, and heat pump integration results

Pinch analysis (PA) was applied to the initial evaporation system design to identify potential opportunities for integrating vapour recompression. For this purpose, the evaporation and condensation heat loads of each evaporator effect were included in the process stream dataset, alongside the background process streams, for use in the PTA. The GCC for the unmodified system indicates that MVR is most appropriately integrated on Effect 1. This configuration has therefore been adopted in the base-case design.

In this configuration, the MVR is considered an integral part of the process rather than part of the heat exchanger network (HEN). Accordingly, the evaporation and condensation loads associated with Effect 1 are removed from the targeting dataset. The integration of the MVR in Effect 1 shifts the pinch temperature to 54.0°C, corresponding to the tube-side saturation temperature of Effect 2.

A second observation from the GCC is that a thermal vapour recompression (TVR) system could be suitable for Effect 2 if external steam utilities were available. However, as the objective of this work is to design a fully electric milk evaporation system, an MVR integrated with Effect 2 is more appropriate. The GCC for the single-MVR configuration (Figure 63 b) provides insights into potential MVR placement strategies for the second effect. Theoretically, partially upgrading the condensing load from Effect 2 would require less work, allowing the remainder of the condensation duty to cascade downstream (Option 1). Alternatively, upgrading most of the condensing load in Effect 2 would maximise the utilisation of the vapour bleed stream from Effect 1 for milk heat treatment, while also enabling better integration opportunities for additional heat pumps (Option 2).

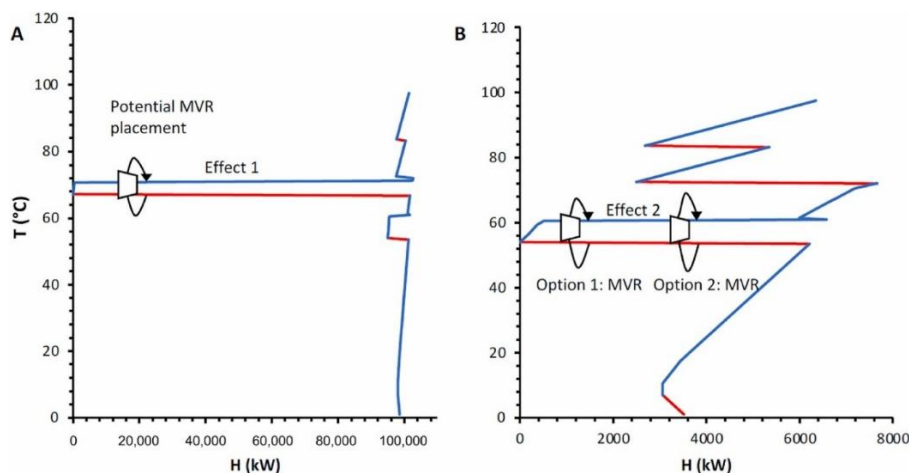


Figure 63 GCC of evaporation system illustrating a design with no vapour recompression (a), and with MVR integrated on Effect 1(b).

Walmsley et al (2016) reached a similar conclusion, recommending MVR integration for both effects. Following this approach, a dual-MVR base-case configuration was developed. The GCC for the MVR/MVR base case presents a revised pinch temperature of 72.5°C, determined by the shell-side pressure of Effect 1 via its vapour bleed stream (Figure 65a). The operating temperature of Effect 1 is already set at the upper acceptable limit to avoid undesirable denaturation of milk proteins, placing a practical constraint on further temperature increases.

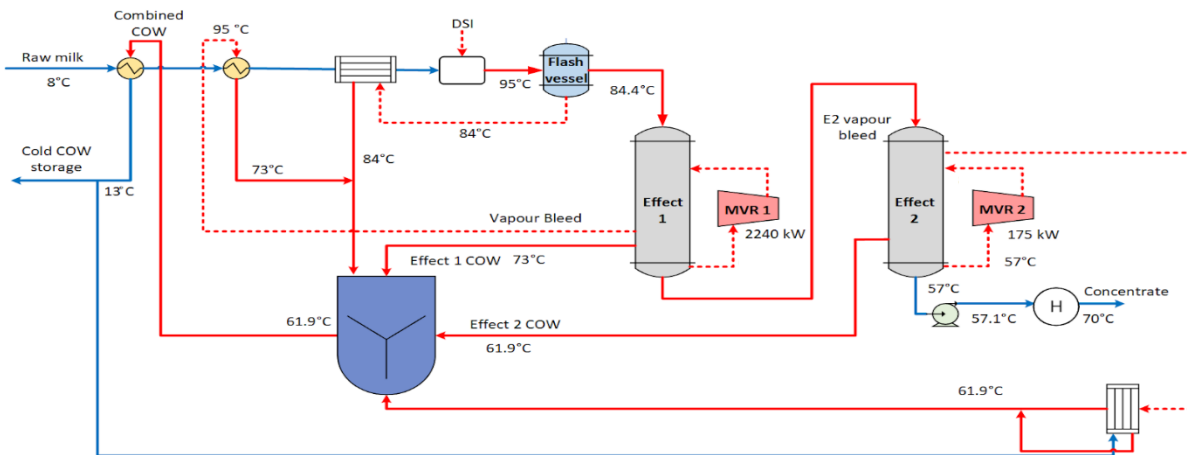


Figure 64 Option 2 a Double MVR System (No Effect 2 Vapour Recycling)




Figure 65 Development of an electric evaporation system (with integrated MVRs on Effect 1 and 2) using the GCC. (a) Non-optimised system design with single DCH. (b) Optimised design with single DCH.

7.4.2 Analysis of pressure set points for flash recovery and effect 2

Analysis of the GCC for the MVR/MVR base case revealed a potential inefficiency associated with the flash vessel upstream of Effect 1 (Figure 65a). To address this, the present study investigated replacing the flash vessel and its associated indirect heat exchanger with a single DCH unit.

DCHs provide rapid heating, which is advantageous in the upper heat treatment temperature range, and are becoming increasingly common in industrial evaporation systems due to their lower capital cost and ease of cleaning (Tyson et al., 2010). In operation, the DCH achieves the required heating duty by flashing the already preheated milk and injecting the generated vapour directly into the incoming cold milk stream. The key thermodynamic distinction between a DCH and a traditional flash vessel with indirect heat exchange is that the DCH introduces additional water into the process milk stream, thereby increasing the downstream evaporation load. This creates a trade-off between improved operability and cleanability offered by DCHs and the lower evaporation load (and associated energy requirement) of traditional indirect heating.

Simulation results indicated that replacing the flash and indirect heat exchanger with a DCH increased power consumption by 55 kW (1.2%) compared to the indirect system. Given the operational advantages, however, DCH units are adopted as the primary configuration for subsequent analyses in this study.

Further optimisation of the system was explored by adjusting the flash pressure set point for the DCH in the Digital Twin. A slight reduction in flash pressure decreased the hot utility target from 4175 kW to 3929 kW (Table 4). At the optimal flash pressure (determined via running an optimisation routine that varied DCH pressure and targeted minimum hot utility), the GCC exhibited a double-pinch characteristic: an upper pinch driven by the flash pressure and a lower pinch associated with excess vapour from the shell side of Effect 1 (Figure 65b).

Table 20 Heat and exergy Pinch targeting for differing cases.

Case	Pinch temperature (s) (°C)	Hot utility target (kW)	Cold utility target (kW)	X _{deficit} (kW)	X _{surplus} (kW)	ΔX _{net} (kW)	Work target, (kW) η _{II} =0.5	Work target, (kW) η _{II} =0.4
A - Non-optimised single DCH	72.5	4175	1138	861	423	438	1511	1983
B – Optimised single DCH	72.5, 81.5	3929	861	808	372	437	1430	1871
C – Non-optimised double DCH	5.5–10.5	3549	456	713	294	419	1279	1665
D – Optimised double DCH	5.5–10.5	3549	456	613	208	405	1122	1449
E – Optimised double DCH (same tube side evaporator pressure)	5.5–10.5	3549	456	613	211	402	1120	1448

Introducing a second DCH recovery unit downstream of the direct steam injection (DSI) stage allows additional heat to be recovered after steam addition, provided that the operating pressures are optimised. This configuration produces a triple-pinch profile on the GCC (Figure 66). The modification reduces the minimum hot utility target from 3929 kW to 3549 kW, representing a further saving of 380 kW. However, the increased heat recovery shifts the process heat profile, creating a larger heat deficit at lower temperature intervals. This occurs because the enhanced recovery reduces the quantity of excess vapour produced by the evaporator effects, thereby decreasing the amount of low-temperature heat available for recovery.

The pinch analysis also revealed an additional potential design benefit. In the baseline DCH configuration (Figure 65), a small plateau appears at approximately 60°C, corresponding to excess vapour from Effect 2, which is generated at a lower temperature due to the lower operating pressure of that effect. Increasing the tube-side pressure of Effect 2 reduces flashing, which in turn increases the mechanical load on its MVR. This adjustment aligns the excess vapour temperature from Effect 2 with that from Effect 1, enabling the two streams to be combined. This simplification benefits the

heat exchanger network (HEN) design by reducing the number of distinct vapour sources that must be integrated.

Furthermore, the availability of additional high-temperature vapour facilitates operation of the associated heat pump (MVR) at a lower ΔT_{lift} , which improves its coefficient of performance (COP). The combination of HEN simplification and improved heat pump efficiency demonstrates how targeted operational adjustments can enhance both thermodynamic performance and process integration in the PI&E framework.




Figure 66 Design with the same tube-side pressure on Effects 1 and 2, and optimised double DCH recovery.

7.4.3 Exergy Analysis for Design Improvement

Exergy analysis is particularly valuable for processes in which both heat and work interactions are significant Hamsani et al. (2018). While one important application is the creation of thermodynamic targets, as outlined in the methodology, its usefulness extends beyond this function. Exergy can also serve as a diagnostic metric to evaluate the potential of a proposed design to deliver energy savings, even when conventional utility targeting suggests little or no change.

In the context of the case study, exergy targets are used to determine whether a process modification could reduce electrical work requirements, even in cases where the hot and cold utility targets remain unchanged. Table 20 presents five representative cases illustrating the advantages of incorporating exergy targeting into a comprehensive PI&E methodology. For this purpose, Equation 58 from the methodology was applied with $\gamma=0$ (full pocket cut), using assumed second-law efficiencies of 40% and 50%, values that are representative of typical heat pump performance (Arpagaus et al., 2018). The full-pocket-cut assumption was selected to provide an approximate but consistent benchmark for comparing alternative designs, rather than to predict absolute achievable performance.

In the single-DCH configurations (Cases A and B), the pinch temperature is initially determined by the shell-side pressure of Effect 1. Following optimisation of the flash pressure, a double-pinch profile emerges, which maximises heat recovery and leads to reductions in both utility and work targets. The reduction in utility demand arises from the interaction between flash pressure and pinch location, consistent with the plus–minus principle. The plus-minus principle states that changes should either increase the duty of hot streams above the Pinch or cold streams below the Pinch (plus) or decrease the duty of cold streams above the Pinch or hot streams below the Pinch (minus)). “Plus” changes will result in reduced hot utility and “minus” changes will result in reduced cold utility. Correspondingly, the work targets predict a reduction in electricity consumption of 81–112 kW due to flash pressure optimisation.

Transitioning to a double-DCH configuration yields substantial reductions in both utility and work targets relative to the single-DCH system. However, for all double-DCH cases (Cases C, D, and E), the hot and cold utility targets are identical and insensitive to changes in either Effect 1 shell-side pressure or flash pressure, owing to the narrow pinch temperature range. In these cases, exergy-based work targets provide the primary mechanism for further optimisation. For example, optimising the flash pressure in Case D reduces the predicted electricity requirement by 157–216 kW compared to Case C. Case E represents a variation of Case D in which the shell-side pressures of Effects 1 and 2 are equalised, simplifying both the design and operation of the HEN.

These results highlight the complementary role of exergy analysis within PI&E, particularly in guiding operational and configuration changes that may not be apparent from utility targeting alone. While the exergy analysis conducted for the DCH placement used the more approximated techniques, it demonstrates the potential of exergy targets to identify meaningful work savings.

7.4.4 *Additional heat pump selection and placement*

Pinch principles can help dictate the placement of heat pumps in a process, however, the technique must be refined when undertaking a full electrification approach, as described in Section 7.2.7. The extended heat pump placement principle can enable lower total electricity consumption for full Process Electrification.

The evaporation process produces GCCs with large pockets of internal heat recovery, which subsequently can be utilised to minimise ΔT_{lift} and hence, provide the greatest work efficiency (Figure 67). Before utilising the pocket, a combined refrigeration heat pump cycle can firstly be integrated around the primary pinch, satisfying the cooling demands fully. Next is, the integration of an air-source heat pump between 25 °C and 50 °C to open the recovery pocket and expose the secondary pinch for further integration.

Refrigerant and cycle selection is a complex topic full with trade-offs. A small refrigerant study was conducted using the high-fidelity model described in Chapter 6. The learnings of this study informed the technology selection for the evaporator. See Appendix Three: Refrigerant Efficiency Studies for the detailed analysis.

From the study the large temperature glide experienced by this heat pump, a trans-critical CO₂ heat pump cycle enables a close match of the temperature profiles within the gas cooler, which subsequently reduces entropy generation. CO₂ as the refrigerant offers acceptable pressure ratios (3–4), which allows the application of high-efficiency compressors. It also is non-flammable and hygienic providing further benefits. A propane heat pump with a sub-cooler is a favourable alternative as well.

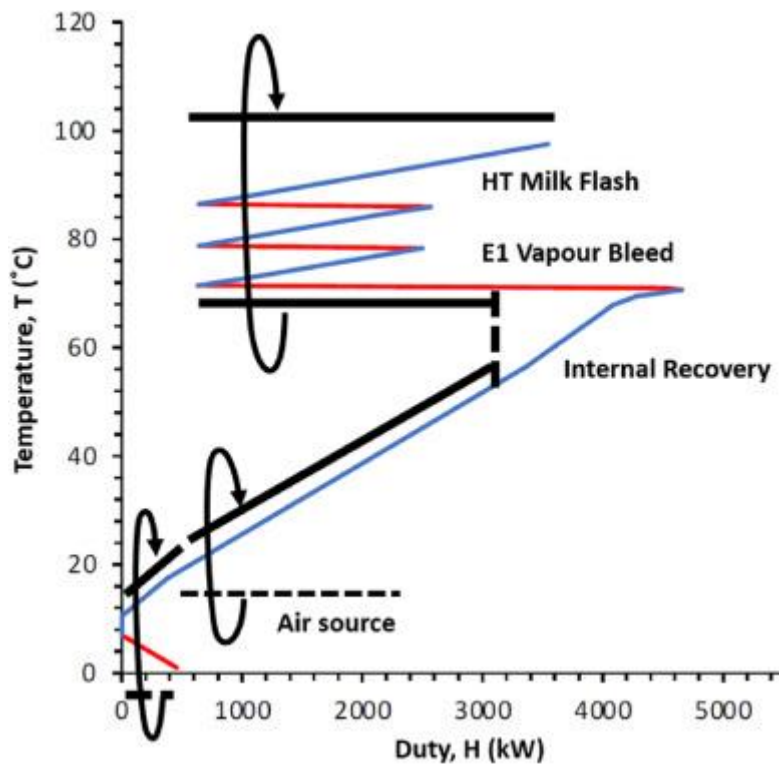


Figure 67 GCC of MVR/MVR base case design with illustration of potential heat pump placement.

After integration of the low-temperature deficit, the Pinch returns to the upper-temperature section of the process, which is labelled the secondary Pinch. Having opened the heat recovery pocket, the excess vapour from the effects can now be upgraded to satisfy the upper heating demands. In this study, three small MVRs were selected and connected in series to provide the 25 °C lift to 98 °C DSI, as it achieved a more favourable COP (10 compared to a high-temperature heat pump COP of 4.91) (Piller, 2025). Also, the MVR did not require an intermediate heat transfer fluid to minimise contamination risk. Additionally MVR is traditionally viewed as a large piece of equipment; however, recent technological advancement has seen MVR fans on the market that require as little as 1.5 m by 1.0 m base dimensions and 1.6 m height for a single-stage (Piller, 2025).

7.4.5 Heat exchanger network optimisation

Each design modification is represented as a grid diagram to improve the selection of the equipment and allow a clear comparison of COPs. Figure 68 illustrates how the grid diagram can be used for efficient identification of key heat exchanger matches and estimation of the required ΔT_{lift} for the integrated heat pumps. The key matches common across the design iterations were Effect 1 and Effect 2 vapour bleed (denoted E1 + E2, COW water, and high-temperature flash vapour streams to the milk).

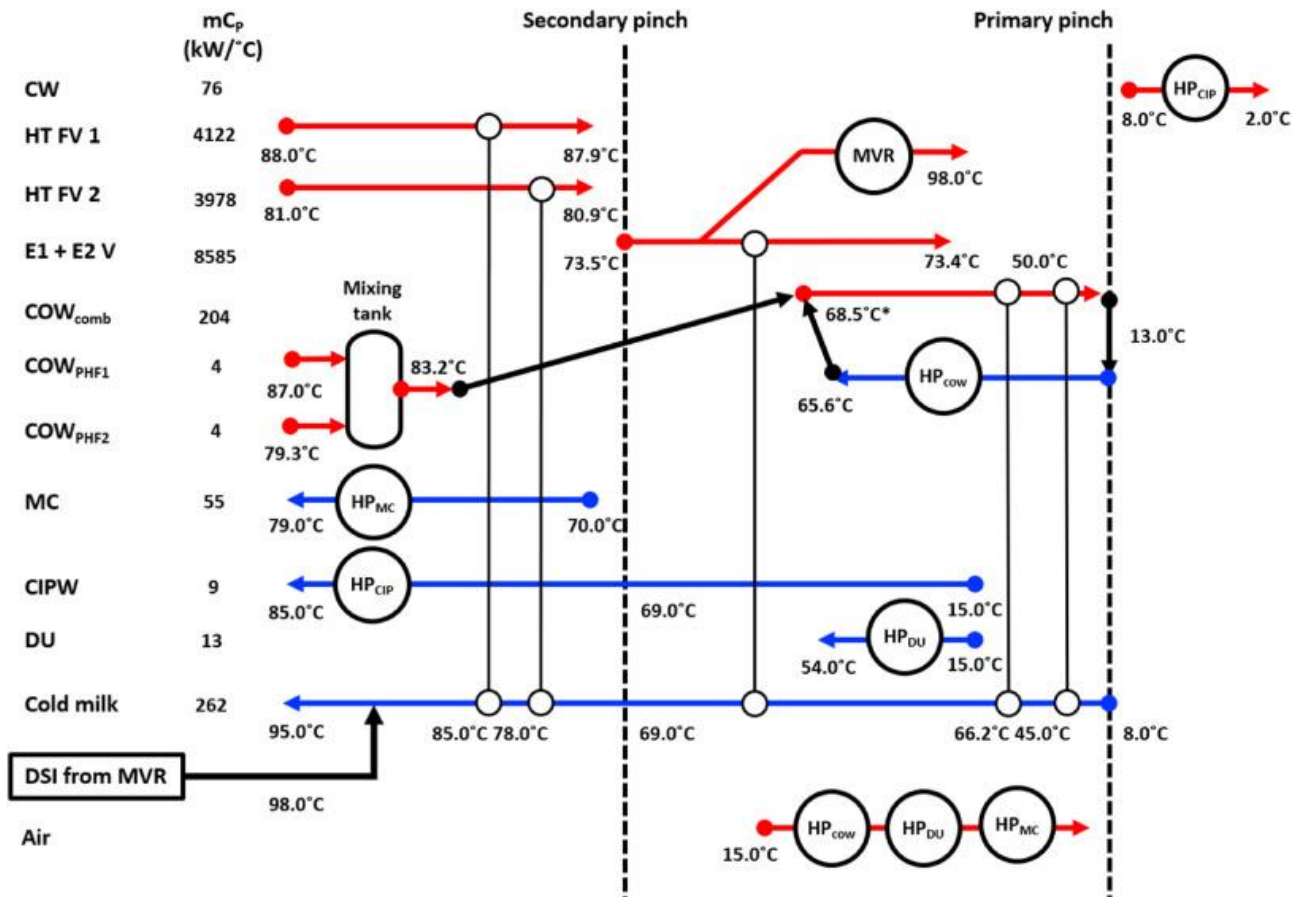


Figure 68 Heat Exchanger Network diagram of the electric milk evaporator.

The grid diagram also helped to visualise the design for practical considerations in the placement of heat pumps. For example, the milk should not exchange heat with a refrigerant directly to prevent leaks contaminating the process milk stream. Using the grid diagram as a guide, a portion of the COW water was recycled and reheated using a CO₂ heat pump from 13 to 68.5 °C, which can then be used for additional heating of the milk stream.

7.4.6 Evaluation of design iterations

After each PI&E analysis iteration, it is necessary to implement and test the changes using the digital twin. As a result, several full process system designs (including the HEN, heat pumps and evaporators) have been simulated and evaluated (Table 21). Comparing process design iterations enabled greater insight into the inherent mechanisms and non-obvious trade-offs present in the problem. Additionally, process designs included simulation of a triple DCH system (instead of a double DCH), a heat treatment plate heat exchanger without flash (i.e., a replacement of the DCH system), and a stand-alone heat pump for the final concentrate preheating.

Table 21 Power requirements for the different design cases.

Options	Design features	DSI required	Power requirements	Specific power requirements	Specific power requirements
1	MVR Effect 1 + MVR Effect 2 + Simple HP	–	4756 kW	159 kWh/t _p	51.3 kWh/t _{w-evap}
2	1 DCH + Integrated concentrate heating	3195 kW	3929 kW	131 kWh/t _p	42.4 kWh/t _{w-evap}
3	2 DCH + Integrated concentrate heating	2648 kW	3535 kW	118 kWh/t _p	38.2 kWh/t _{w-evap}
4	3 DCH + Integrated concentrate heating	2170 kW	3530 kW	118 kWh/t _p	38.1 kWh/t _{w-evap}
5	HX (no DCH) + Integrated concentrate heating	1315 kW	3498 kW	117 kWh/t _p	37.8 kWh/t _{w-evap}
6	HX (no DCH) + Separate HP for concentrate heating	–	3477 kW	116 kWh/t _p	37.5 kWh/t _{w-evap}
7	2DCH + Separate HP for concentrate heating	–	3593 kW	120 kWh/t _p	38.8 kWh/t _{w-evap}

It was the initial thought that greater levels of heat recovery preheat section with a lower DSI requirement would translate to lower total electricity consumption for the evaporation system. Although this general trend was followed, DSI reductions did not always reduce the total required work. For example, the high-thermal efficiency designs (triple DCH and heat treatment plate heat exchanger without flash) reduce the required DSI but overall result in similar total electricity consumption.

The multiple trade-offs present in a full electric evaporation system cause the total electricity consumption to flat line (or even slightly increase) when attempting to minimise the DSI requirement. In the cases above, this “flat lining” resulted from interactions between the heat cascade and the trade-off of differing heat pumps with different COPs. As more heat is recovered in the preheat section, less heat is sent to Effect 1 and, therefore, after the evaporation process in Effect 1, the less excess vapour is bled. This excess vapour is recycled, upgraded, and used for upgrading to DSI by the three-stage MVR in addition to preheating the raw milk. Even though less DSI is required, the decrease in the excess vapour is large enough to cause a larger deficit of heat in the preheat section. As a result, more heat is demanded of the COW recycle heat pump to make up for the larger deficit. The simulation calculated the COW heat pump to achieve a COP of about 4 whereas the three-stage MVR

reached a COP of 10. The lower COP heat pump compensates for the lack of heat cascade by upgrading more air-source heat, nullifying the DSI savings.

7.4.7 Final process design

The PI&E analysis culminated in constructing the final process design (Figure 69). The design begins by heating the milk past 60 °C within one heat exchange to adhere to the set product safety constraints of avoiding rampant thermophilic bacterial growth. This heat exchange would occur in two heat exchangers operating in parallel to avoid disruptions to the operations and for better ease of cleaning maintenance. The arrangement allows for the operating heat exchanger to be shut off when it needs cleaning simultaneously allowing for the second heat exchanger to be activated.

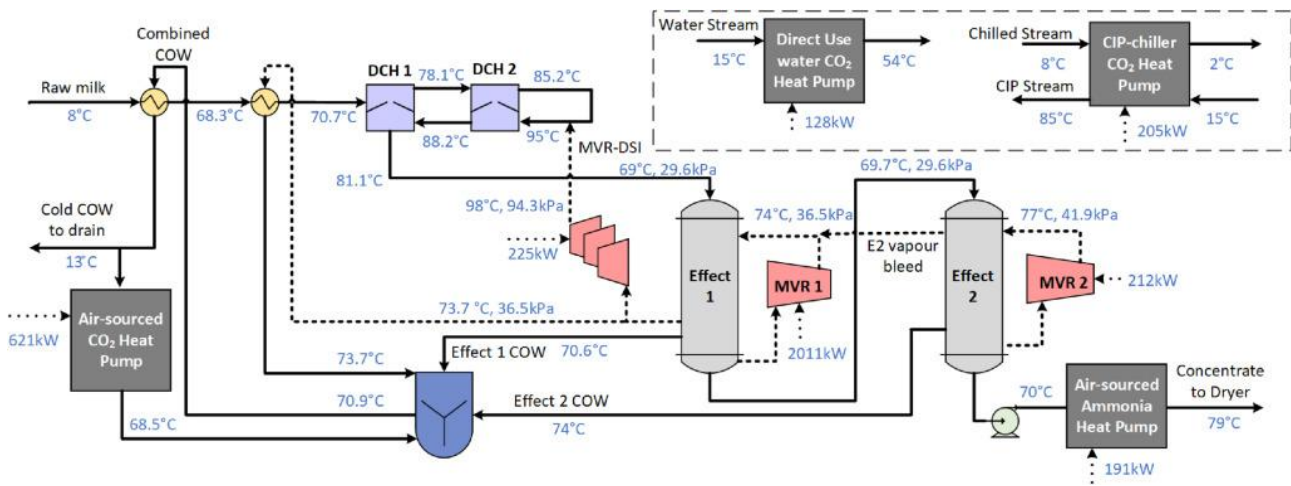


Figure 69 Final process flowsheet design of the fully electric milk evaporation system.

The key features of the preheat section are two direct contact heaters (DCH) and a three-stage MVR. The double DCH system was preferred because it provides fast heat transfer rates, satisfying the production constraints while incurring a minimal electricity consumption penalty. The three-stage MVR utilises approximately 225 kW to upgrade the excess vapour from both effects to substitute steam that would otherwise be provided by a fossil-fuelled boiler.

The design utilises an air-sourced trans-critical CO₂ heat pump, which upgrades the combined COW water from both the effects up to 70.9 °C using 621 kW of electricity. The upgraded COW water is then used to transfer heat to the incoming raw milk up to the specified temperature (68.3 °C). Effect 1 and Effect 2 operate at 69 °C and 69.7 °C respectively (tube-side). MVRs were integrated into both effects to help drive the evaporation process, with the MVR on Effect 1 operating at 2011 kW and

the MVR on Effect 2 operating at 212 kW. The final milk concentrate is then heated indirectly to 79 °C by an air-source ammonia heat pump, which uses 191 kW. The final design achieved a milk solids concentration of 51 wt% which is within the specified design requirement. Separate from the process, there is also an air-sourced trans-critical CO₂ heat pump used to heat direct use water and a dual-purpose chiller and heater for CIP water and chilling demands.

The proposed electric milk evaporator design has a combined COP of 3.17 where a single VHTHP used to provide the heating requirements up to 105 °C would result in a COP of 2.1. Overall, the proposed design requires 3593 kW of electricity, resulting in a 32% operational cost reduction when compared to the previously stated breakeven target of 5262 kW. The proposed design has 82% emissions reduction at 10 kt CO₂-e/year. It is important to note that if the electricity supply to the site is non-renewable these emissions savings will fail to eventuate. It is therefore important to research the grid emissions for the area before electrification.

7.4.8 Comparison with conventional heat pump integration

Bühler et al. (2019) developed several milk powder plant electrification designs that primarily focused on utility system electrification, with comparatively limited redesign of the process system itself. One of these designs, referred to as “Electrification scenario 4”, included an MVR/MVR evaporator system combined with an air-source heat pump. This configuration is broadly comparable to Option 1 in Table 21 and represents a more conventional heat pump integration approach.

To enable a fair comparison, the operating requirements, such as heat treatment temperature, and the boundary conditions, such as initial and final milk solids content, must be kept consistent. Under these equivalent conditions, Option 1 required a total work input of 4756 kW, corresponding to a specific electricity consumption of approximately 159 kWh/tp. In contrast, the recommended electric milk evaporation system developed in this study, Option 7, required 3593 kW, corresponding to approximately 120 kWh/tp. This represents a reduction of 1,163 kW, or approximately 28%, in total work input compared with the simpler single heat pump design.

This comparison demonstrates the effectiveness of the proposed PI&E method. The method is better suited to electrification because it identifies not only where heat is available, but also where it can be upgraded with the lowest avoidable work penalty. The improvement therefore arises from reducing temperature lift, improving heat pump placement, and validating each design iteration within the process simulation rather than relying only on an idealised targeting diagram.

7.4.9 Evaluation of additional operating parameters and sensitivity analysis

Industrial milk evaporators do not all operate at the same temperatures, concentrations, and other set-points. As a result, a successful process design must also be able to operate under a range of conditions. Using the final process design and its digital twin, four key design variables have been varied to analyse the sensitivity of total power consumption, as presented in Table 22. Note the sensitivity values were preprepared in 4.3.5.

Table 22 Sensitivity analysis of key design variables on total power consumption.

Design variable changed	Range examined	Change in total power consumption
Evaporator tube-side temperature	67 °C–72 °C	–16.2 kW/°C
Heat treatment temperature	90 °C–100 °C	+10.6 kW/°C
Final concentration	52%–63%	–13.5 kW/%
Initial concentration (or preconcentration)	14.5%–30.5%	See Figure 70

For the evaporator tube-side and heat treatment temperatures, the relatively small change in power consumption indicates there is flexibility in the operating temperature of the design with minimal energy penalty. To minimise electricity consumption, the evaporator tube saturation pressure and temperature should be at the maximum acceptable temperature to maintain desired product flavour and attributes. Exceeding the maximum acceptable temperature endangers the product quality and may render the product unfit for sale. Similarly, the heat treatment temperature affects product flavour and depends on the specific product and target customer.

Raising the final concentration out of the evaporator system can substantially reduce the energy required in the spray drying operation. However, high-concentrate spray drying is still a developmental technology. This study investigated the effect of producing high solids concentrate (assuming it could be spray dried) on the total work required. The result, in the first instance, appears counter-intuitive where the total work required slightly decreases as the final concentrate solids increase from 52% to 63%. This reduction is the result of a non-trivial trade-off between the various heat pump operations. The simulation shows that MVR Effect 2 would achieve a COP of about 50, due to the small temperature lift required, which means it can efficiently increase the heating load for a low electrical cost. The increase in evaporation results in more hot COW water going to the recovery exchanger, which decreases the heat required from the low COP COW recycle heat pump. In addition,

the low COP concentrate heat pump has less mass flow rate to preheat, which results in a lower heating duty and saves electrical work.

Reverse Osmosis (RO) is an emerging milk pre-concentration technology that is beginning to permeate the dairy industry. As a result, a rudimentary analysis of a milk RO system, as a pre-concentration process, has been added to the simulation to examine the impact of pre-concentrating milk before it enters the primary MVR evaporator system on total power consumption. To date, a clear consensus has not been reached in literature in rating the specific electrical consumption of milk RO. Limited industrial case studies Walmsley et al. (2018) concluded that the typical range of RO costs is between 3.9 kWh/t_{water} to 10 kWh/t_{water} with an average cost of 7 kWh/t_{water} being the general average for milk concentration processes. Each of these specific energy consumption values has been tested to understand the impact, as shown in Figure 70. At the higher energy ratings for RO, there is a slight increase in the total work required, while an energy-efficient RO system could save 21% of the work required. The results look potentially promising for a milk RO pre-concentration system; however, firmer RO energy data is needed.




Figure 70 Effect of changing the initial concentration, because of a pre-concentration reverse osmosis operation, on total power consumption.

7.5 Discussion

The research highlights the significant operating cost reduction that can be achieved through efficient electrification of the evaporation system. There are, however, barriers that need further investigation before enough confidence is generated for general uptake. The current simulations rely on equations for milk component properties from a variety of papers for whole milk. Different milk grades have

different properties, which means the simulations would need re-evaluation to confirm their applicability. A second barrier is the prohibitive investment cost to construct a pilot-scale or full-scale electrified test plant to verify the simulation results. As a result, a collaboration of processing site owners, original equipment manufacturers, and the government is needed to identify the ideal ways to reliably verify the results. A third barrier is the difficulty of estimating and optimising plant capital costs. Current industrial capital costs are highly uncertain due to the global market and ongoing inflation challenges. As a result, this study has focused on the operating cost savings, while also recognising the need to include detailed capital cost estimates in the PI&E design method. An area of opportunity is the retrofit of existing sites. This study has not investigated the retrofit challenge; however, the PI&E design method is adapted to develop retrofit solutions in Chapter 8.

The study considered the milk evaporator system as an isolated process; however, the evaporator plant is a process step in the drying of milk powder and therefore part of a wider system. A total site Pinch analysis would likely highlight additional heat sources and integration opportunities that could further reduce costs and emissions.

7.6 Conclusion

This Chapter presented a novel design for a highly efficient, fully electric milk evaporation system developed through an effective Process Integration and Electrification, PI&E, design method. Encompassed in the PI&E design method was an iterative process that utilised a design digital twin, Pinch Analysis, heat pump placement and selection, and exergy analysis. This study takes a step further than existing literature by using practical considerations to underpin the design decisions over solely minimising energy consumption, reducing barriers to uptake in this respect.

Seven design options were developed and analysed through the proposed method, using a two-effect evaporator with Mechanical Vapour Recompression (MVR) on each effect as a commonality across all designs. The final design included an innovative combination of a three-stage MVR fan, that leverages new compact MVR technology, and two direct contact heaters for milk heat treatment through direct steam injection before evaporation. Additionally, transcritical CO₂ heat pumps and an air-source ammonia heat pump were used to supply heating and cooling requirements within and external to the process.

The work requirement of the final design was 3593 kW (120 kWh/t_p) with an overall COP for supplying heat to the process of 3.17. Compared to the base case design, the work requirement of the

final design represents a 32% operational cost reduction and a 82% emissions reduction. While the final design (Option 7) requires 64 kW (4 kWh/t_p) more work than an equivalent design that uses a heat exchanger instead of direct contact heating for preheating (Option 6), the fast-heating rates and ease of operability of using direct contact heating far outweighed the comparatively small increase in work requirement.

A sensitivity analysis found that the final design remains operable over a variety of operating conditions, including evaporator tube-side pressure, heat treatment temperature, and final product concentration with a negligible increase in total power consumption (<1% per °C or wt %). The sensitivity analysis also considered, at a high level, the potential use of reverse osmosis with the new design. It was found further work savings of up to 21% could be achieved with an efficient RO system presenting a promising area for further research.

Chapter 8 Process Integration and Electrification for Retrofit Design

8.1 Introduction

While the previous chapter focused on greenfield process design, where electrification and integration strategies can be optimised from first principles, the reality of industrial decarbonisation is that the majority of systems already exist and must be retrofitted. Retrofitting presents a very different set of challenges compared to greenfield design. In new-build plants, unit operations, heat exchanger networks, and utility systems can be designed holistically to minimise energy consumption and integrate electrification technologies from the outset. By contrast, retrofits are constrained by the physical, operational, and economic characteristics of the existing site.

Key challenges include the limited flexibility of existing heat exchanger networks, equipment designed around steam-based utilities rather than electrical technologies, spatial and layout restrictions, and the need to maintain product quality and throughput during and after modifications. In addition, retrofit solutions must contend with shorter payback expectations and tighter capital constraints, as the investment competes with the ongoing operation of depreciated assets.

Despite these challenges, retrofit strategies offer significant potential for decarbonisation. By leveraging exergy and process integration insights, it is possible to identify opportunities where targeted electrification interventions can achieve meaningful reductions in energy use and emissions without requiring wholesale system replacement.

This chapter therefore shifts the focus from the unconstrained optimisation of Chapter 7 to the practical realities of integrating electrification into existing milk evaporation systems. The methodologies (with key adjustments) developed earlier are applied here to guide retrofit strategies that balance thermodynamic efficiency with practical feasibility.

Therefore, the concluding addition to the thesis aims to develop a comprehensive PI&E retrofit method that includes potential modification of the process (e.g., operating set-point, change in technology), HEN (e.g., heat transfer enhancement), and process electrification (e.g., MVR). The method represents an extension of Bridge Analysis (Bonhivers et al., 2017), heat pump Bridge Analysis (Schlosser et al., 2021) and the Modified Energy Transfer Diagram (METD) (M.R.W.

Walmsley et al., 2017). The case studies for this work include three typical milk evaporator designs that are found in the New Zealand dairy industry.

The novel contributions of this work to the literature include:

- The development of an improved PI&E retrofit method using an extension of heat pump bridge analysis to include the process units. The inclusion of process unit heat flows allows for process modifications to be considered. A key aspect identified in previous literature to electrify processes efficiently.
- The demonstration of applying the method to multiple related case studies
- The design of common retrofit solutions for milk evaporator plants to transition to renewable electricity fuel.

8.2 Method

The purpose of this study is to produce an intuitive PI&E retrofit method that will be used to determine general cost-effective electrification retrofit strategies for the milk evaporation process. Figure 72 presents a flowchart of the method with descriptions of the key steps described below. Note that some of these steps align with the methods discussed in previous Chapters and therefore are only described in an outline.

8.2.1 Preparation steps

- 1. Gather process design data for the existing process (Chapter 4).
- 2. Define PE targets for full or partial electrification. The targets could also be defined as a percentage of steam supplied by fossil fuels, referred to as fossil steam, to be replaced. The targets should consider practical aspects of the process, such as scheduling, process requirements and physical dimensions of the equipment. The targets should also consider existing electric technologies and operational costs (Chapter 4).
- 3. Create a digital twin for the existing process to generate a reliable and complete data set (Chapter 5 and 6).

8.2.2 Identification of potential retrofit opportunities

The identification of retrofit opportunities begins with a detailed analysis of the existing heat exchanger network (HEN). Conventional pinch analysis typically treats the process as a whole, aggregating all hot and cold streams into composite curves to determine minimum utility requirements. While this approach is useful for new designs, it is less effective for retrofit applications where the existing HEN is fixed, and modifications must be targeted at specific exchangers. To address this limitation, the Exchanger Grand Composite Curve (EGCC) is applied, which is constructed from the Exchanger Problem Table Algorithm (E-PTA) (Lal et al., 2018).

Unlike the standard PTA, which aggregates all hot and cold streams into composite curves, the E-PTA treats each exchanger as a stand-alone problem. For a given exchanger, the hot and cold streams are first shifted by their minimum approach temperature contributions, producing a shifted composite curve (SCC). From this SCC, the heat cascade is constructed across defined temperature intervals, using the following relationships:

$$\Delta T_{i,i+1}^* = T_i^* - T_{i+1}^* \quad \text{Equation 60}$$

$$\Delta CP_{i,i+1} = (CP_h - CP_c)_{i,i+1} \quad \text{Equation 61}$$

$$\Delta \dot{H}_{i,i+1} = \Delta T_{i,i+1}^* \cdot \Delta CP_{i,i+1} \quad \text{Equation 62}$$

The cumulative enthalpy cascade across the exchanger is then calculated as:

$$\dot{H}_i = \dot{H}_{i-1} + \Delta \dot{H}_{i-1,i} \quad \text{Equation 63}$$

Boundary conditions depend on exchanger type: for recovery exchangers, the first and last cascade values must be zero, whereas for heaters and coolers they correspond to utility duties (HU or CU, respectively). Plotting the cascade H against shifted temperature T^* yields the EGCC, which reveals temperature intervals where deficits or surpluses exist, as well as proximity to approach constraints.

To extend this analysis across the full heat exchanger network, the individual EGCC results are combined into a Heat Surplus–Deficit Table (HSDT), which provides a structured overview of energy imbalances across all temperature intervals.

For each interval, the net enthalpy balance is calculated as:

$$\Delta\dot{H}_{i,i+1}^{(HEN)} = \sum_k \Delta\dot{H}_{i,i+1}^{(k)} \quad \text{Equation 64}$$

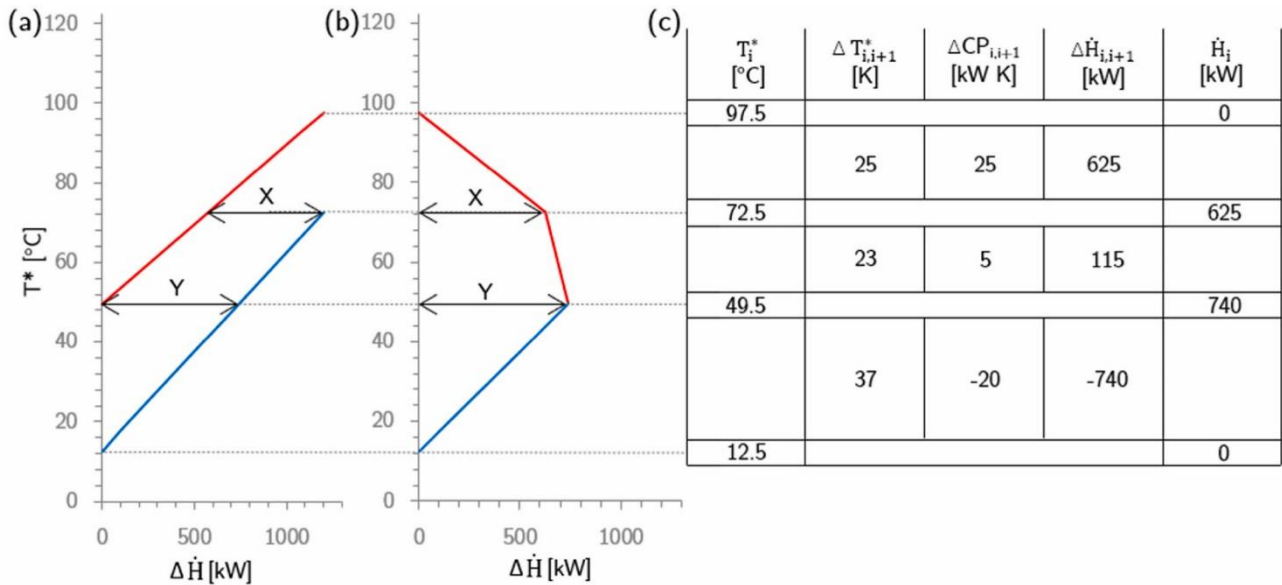


Figure 71 The relationship between the Exchanger Shifted (a) Composite Curve (ESCC), (b) EGCC, and (c) E-PTA for recovery exchanger as shown in Walmsley et al (2017)

In this tabular representation, positive values denote temperature intervals where heat surpluses exist, while negative values indicate deficits. The HSDT acts as a numerical counterpart to graphical methods such as the Modified Energy–Temperature Diagram (METD), but with added detail on the shifting potential of individual streams. By explicitly displaying surpluses and deficits at each temperature interval, the HSDT highlights cross-pinch violations and areas where additional integration (through heat recovery or heat pump bridging) could reduce utility consumption.

Together, the EGCC and HSDT provide a rigorous framework for diagnosing inefficiencies in the current heat exchanger network and for prioritising retrofit interventions. They allow engineers to determine whether retrofitting an existing exchanger match, shifting operating conditions, or introducing electrification technologies such as heat pumps or vapour recompression would provide the greatest benefit in reducing utility demand.

	C2	C1	E1	E2	H1
T*	ΔH_{net}	ΔH_{net}	ΔH_{net}	ΔH_{net}	ΔH_{net}
°C	kW	kW	kW	kW	kW
245					
	-	-	-	150	-
235					
	-	-	-	600	-1,200
195					
	-	-	83	50	-100
191.7					
	-	100	167		-200
185					
	-	600	1,000	-800	-1,200
145					
	-	690	230	-	-
99					
	600	360	-480	-	-
75					
	-	600	-800	-	-
35					
	-	-	-200	-	-
25					
$\Sigma \Delta H_{Surplus}$		2,350	1,250		
$\Sigma \Delta H_{Deficit}$			-1,480		-1,500

Equation 65 HSDT illustration from Lal et al (2018)

8.2.3 Expansion of the HSDT Method to include processes

Conventional Heat Surplus–Deficit Tables (HSDTs) are typically applied to the heat exchanger network (HEN) and are primarily used to identify retrofit opportunities through exchanger modification or re-matching. However, the scope of analysis can be expanded to include process operations themselves. This extension provides a broader perspective, enabling the identification of additional opportunities for emissions reduction and process electrification beyond the HEN alone.

To include process operations, the first step is to define a system boundary around the unit operation of interest. Unlike heat exchangers, which involve straightforward heat transfer between streams, process units may include reactions, separations, or mixing, making their thermal behaviour more complex. In these cases, conventional methods of defining stream data as enthalpy differences between supply and target conditions cannot be directly applied.

Instead, the approach requires the definition of a reference temperature that acts as an intermediate or imaginary state. All inlet streams to the process are “cooled” to the reference temperature, while all outlet streams are “heated” from this reference temperature back to their actual supply conditions. By bringing all streams to a common reference point, the process unit can be represented in a temperature–enthalpy diagram as an enclosed “pocket.”

An enclosed pocket signifies that the surpluses and deficits of the process are internally balanced, but the shape and position of the pocket provide valuable information about integration potential. For example, a pocket may indicate opportunities for coupling the process to surrounding streams or electrification technologies such as heat pumps. Importantly, this approach maintains consistency with conventional analysis while extending its applicability to more complex operations.

By incorporating process units into the HSDT framework, retrofit studies can account for both the HEN and internal process balances, thereby creating a more holistic understanding of where energy recovery and process electrification can be most effectively implemented.

8.2.4 Identify waste heat to “unlock” higher temperatures

The identification of waste heat is carried out by processing the stream data into a Heat Surplus–Deficit Table (HSDT) format. In addition to the conventional format, extended versions of the HSDT can incorporate explicit rows for utilities and potential heat pump integration, allowing rapid appraisal of electrification opportunities.

A potential heat pump bridge is defined as the integration of underutilised heat otherwise rejected to cooling utilities or used in high-emission operations into useful process heating duties via a combination of recovery exchangers and heat pump devices. Waste heat sources are therefore identified by two key indicators: (1) heat surpluses directly connected to cooling utilities, and (2) heat surpluses currently serving processes reliant on fossil-based utilities.

A consideration in evaluating waste heat recovery is the temperature lift required by the heat pump. A smaller temperature lift results in higher COPs, thereby reducing the associated electricity consumption. To assess this, Bridge Analysis is employed to establish feasible utility paths that connect waste heat surpluses to process heat demands at higher temperature intervals. Each path is tested against the required minimum temperature approach, which ensures realistic integration.

Because HSDTs are generally presented on a non-shifted temperature scale, the minimum approach temperature constraint must be manually applied during the analysis.

The outcome of this procedure is a set of potential waste heat recovery opportunities, expressed as feasible bridges between low-grade surpluses and higher-grade demands. These bridges form the basis for subsequent evaluation of process electrification strategies, including the selection of appropriate heat pump technologies.

8.2.5 Upgrading waste heat

Once waste heat sources have been identified, the next step is to evaluate options for upgrading this heat to useful process temperatures through electrification technologies, most commonly heat pumps. The selection and integration of heat pumps must take into account process requirements and conditions, including whether the heating duty requires direct heating (e.g., contact with process streams) or indirect heating (via heat exchangers). These practical considerations determine the most appropriate technology for the retrofit solution. For example, direct heating duties can be well-suited to open-cycle systems such as Mechanical Vapour Recompression (MVR), while indirect heating may favour closed-cycle heat pumps.

The performance of a proposed heat pump integration is typically assessed using the COP, which provides a measure of the ratio of useful heat delivered to electrical energy consumed. COP values can be estimated in three ways:

1. Correlations based on the mean temperature lift of the heat pump, or
2. Efficiency models, such as applying an isentropic efficiency to MVR systems.
3. Digital Twin of the Heat Pump System

An essential consideration is the temperature lift (ΔT_{lift}) between the source (waste heat) and the sink (process demand). Smaller temperature lifts generally result in higher COPs, thereby reducing electricity consumption and improving economic feasibility. Conversely, larger temperature lifts substantially reduce COP, increasing operational costs and reducing competitiveness of the retrofit.

8.2.6 Calculate levelised cost of heating

The Levelised Cost of Heating (LCOH) provides a comprehensive economic measure of heating supply technologies by integrating capital, demand-related, and operational costs over the equipment lifetime. For retrofit applications, it enables direct comparison between electrified heating technologies (e.g., MVR, heat pumps) and conventional fossil-fuel boilers on a consistent cost-per-unit-heat basis.

The LCOH was calculated by summarising the capital, demand-related, and operation-related costs according to Equation 66, Equation 67. The methodology accounts for the useful heat supplied, steam savings, compressor capacity, and relevant economic parameters.

$$LCOH_i = \frac{A_{cap} + A_{dem} + A_{op}}{Q_{sav}} = \quad \text{Equation 66}$$

$$\frac{a \cdot f_{c,inst} \cdot c_{i,s} \cdot P_{com} + p_{ref} \cdot r_p \cdot t_{FL} \cdot P_{com} + f_{M,c} \cdot c_{i,s} \cdot P_{com}}{Q_{sav} \cdot t_{FL}}$$

$$a = \frac{(1 + i)^N \cdot i}{(1 + i)^N - 1} \quad \text{Equation 67}$$

Capital costs were annualised using the annuity factor, which spreads the investment over the equipment's economic lifetime. These costs depend on the annuity factor, the specific investment costs, the run index, and the installed compressor power. The annuity factor itself is a fitting parameter calculated according to Equation 67.

Demand-related costs were calculated from the reference energy carrier price, the ratio of electricity price to the reference fuel price, and the annual operating hours at full load. This component reflects the cost of energy consumption required by the electrified system relative to the conventional baseline.

Operational costs were determined from a maintenance factor applied to the run index, investment costs, and compressor power. This ensures that fixed and variable maintenance expenses are accounted for in proportion to both capital intensity and energy use.

Table 23 Economic Parameters for LCOH calculation

Symbol	Definition	Unit
$LCOH_i$	Levelised cost of heating for run index i	\$/kWh _{th}
A_{cap}	Annualised capital costs	\$
A_{dem}	Demand-related costs (energy consumption)	\$
A_{op}	Operation-related (maintenance) costs	\$
\dot{Q}_{sav}	Useful heat supplied / steam savings	kW
a	Annuity factor (Eq. 3)	–
$F_{c,inst}$	Installation factor for capital costs	–
$c_{i,s}$	Specific investment cost for run index i	\$/kW
P_{com}	Compressor capacity (installed electrical power)	kW
P_{ref}	Reference energy carrier price (e.g., natural gas or steam)	\$/kWh
R_p	Electricity-to-reference energy price ratio	–
t_{fl}	Operating time at full load	h
$F_{M,c}$	Maintenance factor	–
i	Discount rate	–
N	Economic lifetime	years

8.2.7 Iterative retrofit evaluation

The final stage of the retrofit methodology involves modifying the process simulation to incorporate the proposed retrofit solution. The modified system is then re-evaluated to determine whether the design achieves the desired level of PE and emissions reduction.

If the target PE level is not met, the methodology proceeds iteratively: additional retrofit options are identified, assessed, and integrated into the Twin. Each iteration provides updated performance indicators, including utility savings, and LCOH, which are compared against the defined objectives. This iterative approach ensures that retrofit solutions are systematically refined and that synergies between multiple modifications are captured.

Beyond achieving technical feasibility, the iterative framework also provides a structured pathway to develop staged implementation plans. If full electrification is not immediately viable due to economic or practical constraints, earlier iterations can serve as intermediate retrofit solutions, offering incremental improvements in energy efficiency and emissions reduction while maintaining flexibility for future upgrades.

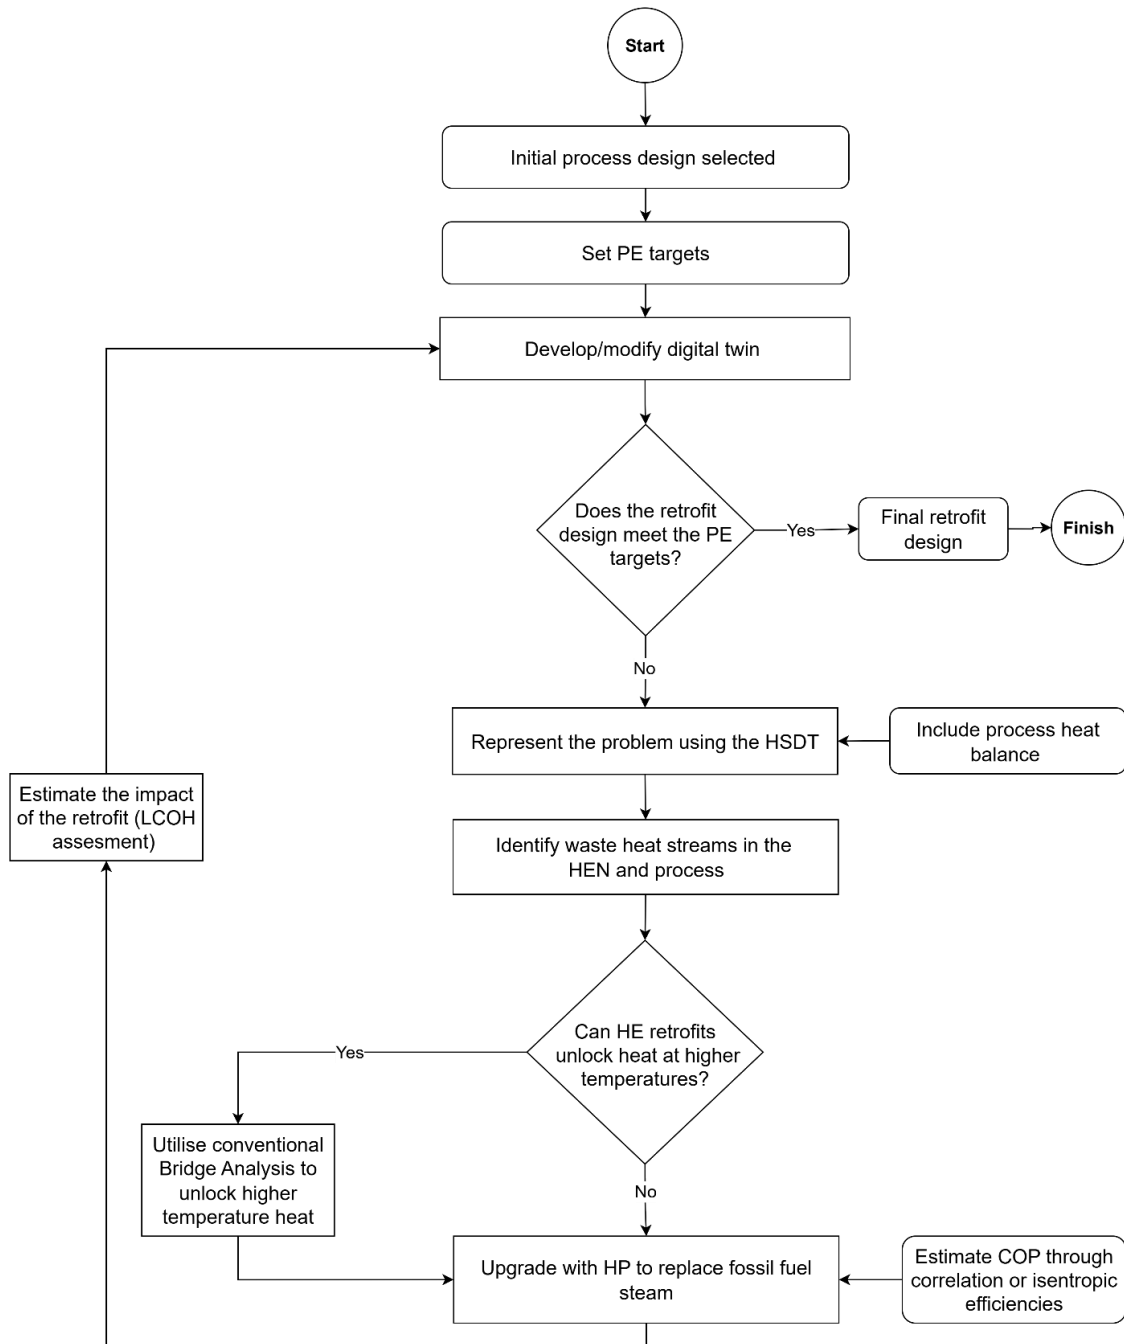


Figure 72 Process Integration and Electrification retrofit design method through iterative and incremental changes.

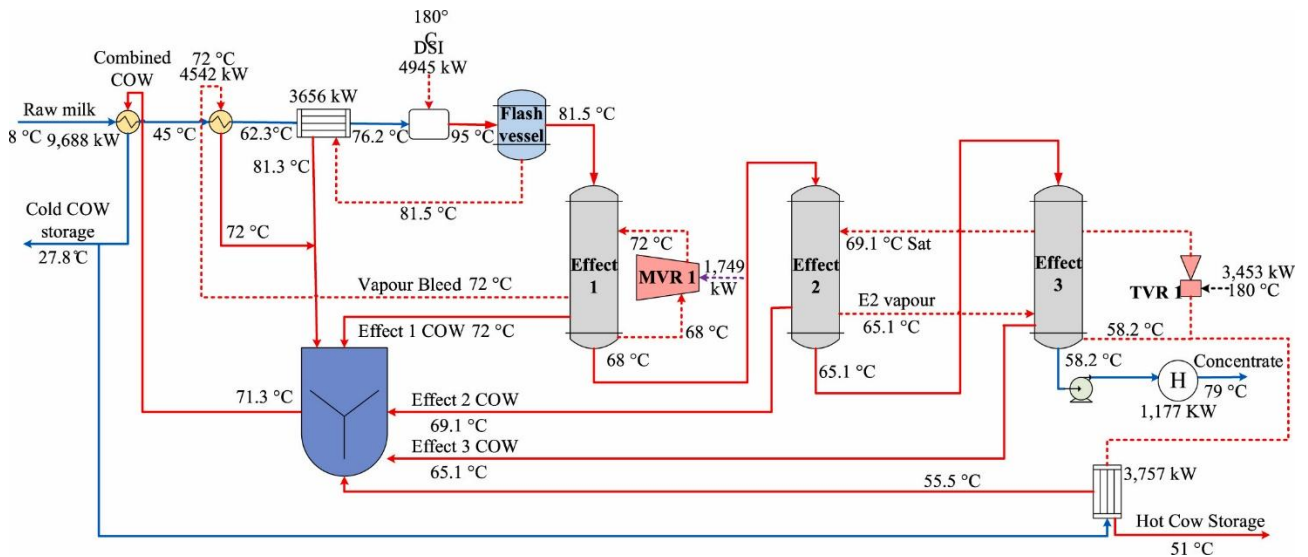


Figure 74 Case 2 – three effect MVR-TVRE evaporation system.

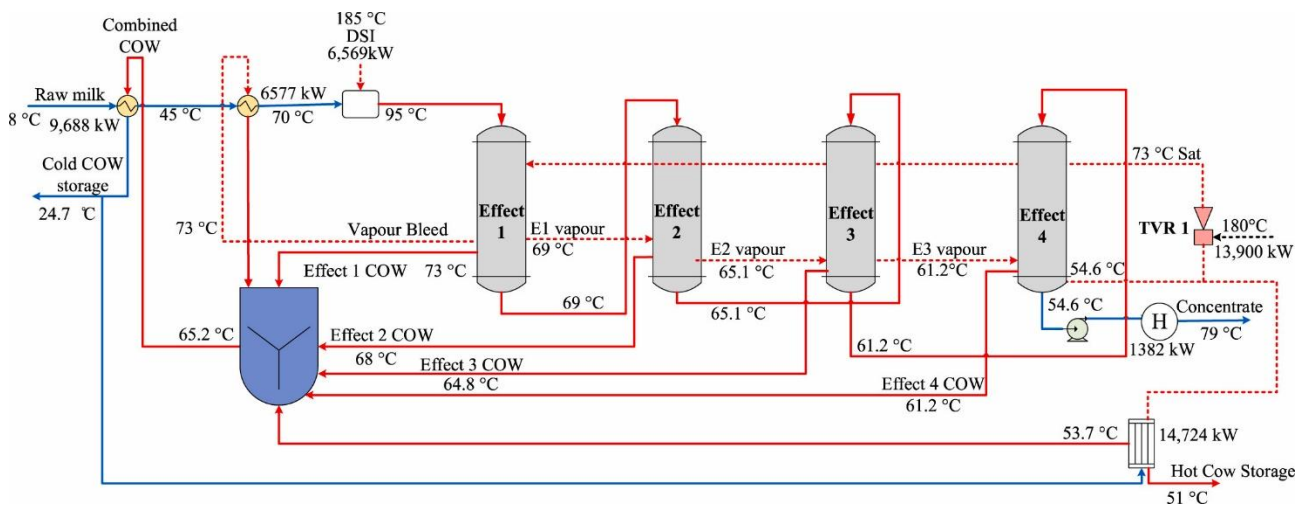


Figure 75 Case 3 – four effect TVRE evaporation system.

The boundary conditions for the evaporation system are the same across the three designs and define the temperatures, compositions, and mass flowrates at the inlet to the preheating section, the inlet to the first effect and the outlet of the evaporation system. The inlet conditions to the preheating section are 247 t/h of whole milk with 14.5 wt% solids entering at a temperature of 8 °C. The milk is preheated to 95 °C with the final heating achieved through direct steam injection for rapid heating and held for about 30 s for heat treatment, before it enters the first evaporator effect. The final concentrate conditions exiting the evaporator are 51 wt% solids and at a temperature of 79 °C.

In Case 1 (Figure 73) and Case 2 (Figure 74), a flash vessel for heat recovery is included between the heat treatment process and the first effect. The flash process recovers 3152 kW. Across the first

effect a mean temperature driving force of 4 °C is maintained for both Cases 1 and 2 using an MVR fan, whereas the final TVR effects requires a driving force of 7 °C due the increase in viscosity of the milk. Case 2 has an additional effect that receives upgraded vapour from the TVR on the final effect. The additional effect then cascades vapour to the final effect. Case 3 (Figure 75) has four cascade effects where vapour exiting the final effect is upgraded using a steam ejector, or TVR unit, and medium-pressure steam to supply the first effect. All effects operate at a mean 4 °C temperature lift except the last which operates at a 7 °C lift.

In this study, the evaporation plant is considered in isolation from the rest of the plant, however, in future research, the consideration of the entire plant when integrating may lead to further electrification strategies. The data for the case studies was generated through process simulation, previously published articles, and industrial experience of the authors in analysing dairy processes.

8.4 Results and discussion

The data generated from the digital twin is then validated against the process design data to determine whether the simulation is representative of the existing process. This is necessary to ensure that retrofit solutions developed in simulation are applicable to the real process. The validated data is then processed using the OpenPinch Excel Workbook to generate the HSDT and undertake the retrofit analysis. To demonstrate the method, the results for Case 1 are presented in a detailed step-by-step manner. Throughout each step, a discussion of the method and the subsequent retrofit design concepts are presented. The same method was also applied to Cases 2 and 3.

8.4.1 Simulation and expansion of the HSDT method

Conventionally, HSDTs typically focus on the HEN and is used to search for heat exchanger retrofit opportunities. However, expanding the analysis to include process operations creates the opportunity to identify further emissions reduction and PE potential. The study presents a modification to stream data collection to allow for these internal balances of process operations. First, a system boundary is drawn around the process to be included (in this case the evaporator (Figure 76)). Unlike a heat exchanger, process operations usually have reactions, separation and/or mixing of streams occurring within the unit. As a result, traditional techniques for representing stream data as enthalpy differences between supply and target conditions cannot be applied. Likewise, the concept of a minimum approach temperature also cannot apply since mixing occurs. Instead, a reference temperature must be defined for the system, which is used as an (imaginary) intermediate state, where inlet streams are

“cooled” to the reference state (target) and outlet streams are “heated” from the reference state (supply). This concept can be applied to a heat exchanger to achieve the same result as conventional analysis. By bringing all streams to a common reference temperature, process units can be mapped on a temperature-enthalpy plot as enclosed pockets, similar to a heat exchanger. An enclosed pocket indicates that the heat surpluses and deficits around the operation is balanced. An example set of stream data for a milk evaporator is presented in Table 24.




Figure 76 System boundary analysis of the evaporator (Case One)

Table 24 Stream data for the system boundary analysis of the evaporator (Case One)

Stream	T_s (°C)	T_t (°C)	ΔH (kW)	ΔT_{cont} (°C)
Effect 2 Entry	57	5	4,267.2	0
Milk Conc 2	5	57	2,806.6	0
TVR IPS	61.9	61.8	6,310.9	0
TVR IPS Sub Cooled	61.8	5	616.4	0
Water Vap	56.9	57	7,125.2	0
Water Vap Sub Cooled	5	56.9	644.5	0
COW Condensate	5	61.8	618.1	0

8.4.2 Identify waste heat and application of conventional bridge analysis

The resulting stream data is processed using the method proposed by Walmsley et al. (2017) and converted into the HSDT format. The study presents a new format style for the table in Figure 77,

adding rows for utility usage and heat pump analysis. The new format explicitly quantifies the degree of PE on the diagram, allowing for quick appraisal of potential retrofit solutions.

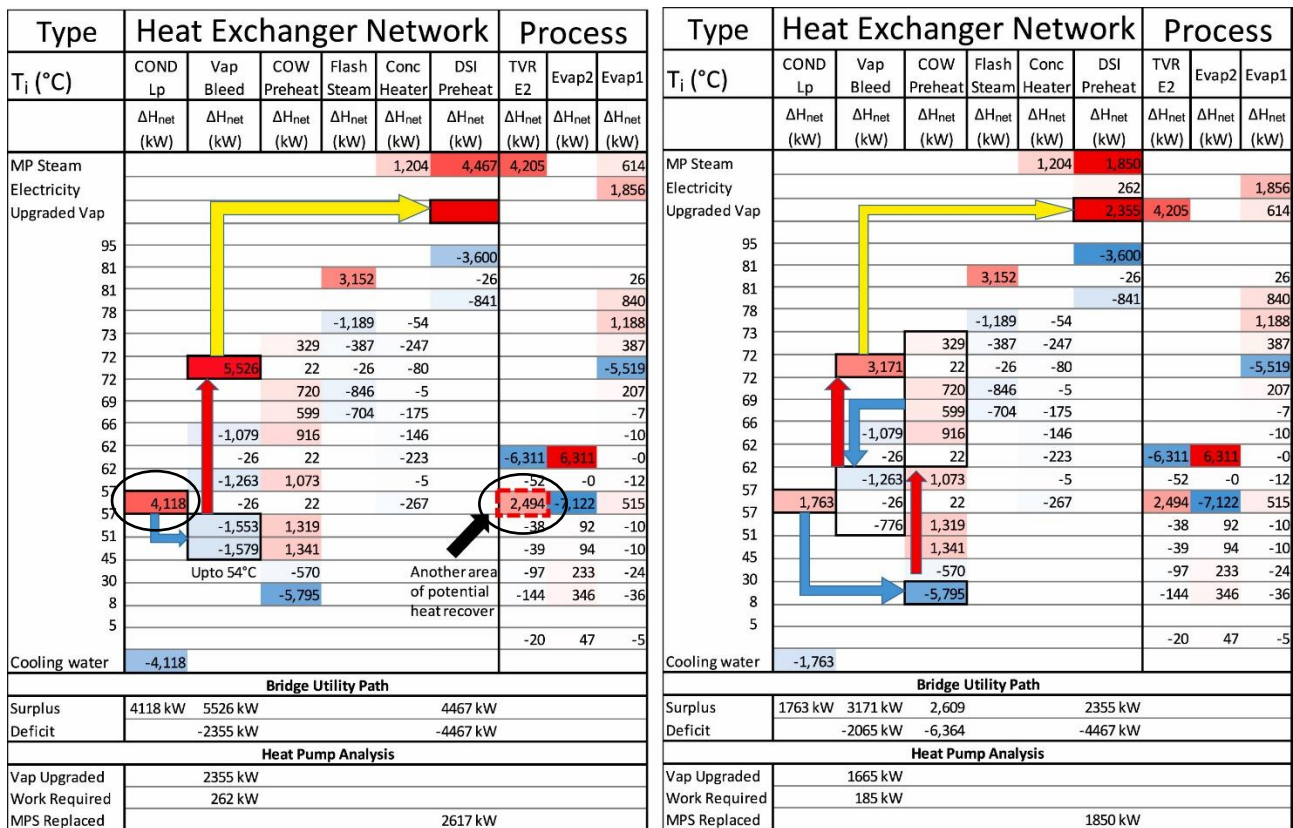


Figure 77 Example Heat Pump Bridge Analysis for the two effect MVR-TVRC Case.

A potential heat pump bridge is formed by integrating a portion of underutilised excess heat through a network of recovery exchangers and heat pump devices to replace fossil steam. Areas of underutilised heat (or areas of recovery potential) are identified through one of two characteristics. 1) The heat surplus is connected to a cooling utility or 2) the heat surplus is utilised in a high-emitting process operation (such as a TVR). Figure 77 identified two areas of underutilised heat, the first being the 4118 kW of surplus cooled in the COND L_p stream. The second is the 2,494 kW of heat surplus upgraded in the TVR E2 stream (highlighted in the circled cells).

A key step in PI&E Retrofit is the minimisation of the temperature lift performed by the heat pump. A smaller temperature lift correlates to a higher efficiency/COPs, minimising the required electricity consumption. Using conventional Bridge Analysis (Lal et al., 2018), a utility path was created from areas of underutilised heat, through a set of recovery exchangers, to a heat surplus at a higher temperature interval. The paper used a minimum approach temperature of 3 °C, however, due to the

HSDTs having non-shifted temperatures, the minimum approach temperature must be enforced by the user.

Figure 77 illustrates how the Heat Pump Bridge Analysis was used to identify retrofit opportunities by first locating available low-temperature heat sources and then determining whether this heat could unlock higher-temperature recovery elsewhere in the process.

In the first case, the identified source stream, LCOND, is available at approximately 57 °C. Two cold streams, the vapour bleed exchanger and the COW preheating exchanger, require heat below this temperature and are therefore feasible initial matches. Transferring heat from LCOND to these cold streams reduces their existing heating demand and, more importantly, releases heat from the corresponding hot sides of these exchangers at higher temperatures. This “unlocked” heat can then be considered for upgrading.

As shown in the first pathway in Figure 77, the heat released through this match occurs at a temperature level suitable for upgrading using a heat pump, allowing displacement of DSI steam. The second pathway follows the same logic but introduces an additional intermediate recovery step, where heat is first matched between the COW preheating stream and the vapour bleed stream before the remaining higher-temperature heat is upgraded. In both cases, the figure should be read as a sequence of heat-recovery and heat-upgrading steps, rather than as a single direct heat pump match.

The first match in Figure 77 shows that only 2,335 kW out of the 3,132 kW is at an approach temperature of 3 °C. Figure 77 also shows a complete 2,355 kW heat bridge from the 57 °C COND L_p stream to the 72 °C Vap Bleed heat surplus.

8.4.3 Integration of heat pump technology

It is important during the integration of heat pumps to consider the process requirements and conditions, such as direct or indirect heating. These practical considerations will inform the type of heat pump technology best suited for the proposed retrofit solution. For the example above, integrating a vapour stream to replace a DSI is a form of direct heating which is suitable for the use of MVR. MVR provides the additional benefit of no intermediate refrigerant and, therefore, higher efficiencies can be achieved. To estimate the electricity cost of a potential bridge, the heat pump COP was calculated. This can be done through heat pump COP correlations (Schlosser et al., 2020a) or isentropic efficiencies of MVRs. For the example problem, an isentropic efficiency of 85 % was

used based on available MVR technology (Piller, 2022). The resulting COP for the 72 °C–98 °C temperature lift was calculated to be 10. It will cost 262 kW to upgrade the 2,355 kW of heat unlocked by the first retrofit bridge.

To demonstrate the impact of smart heat pump integration in economic design, another retrofit option was considered. If the retrofit choice was to instead upgrade the excess vapour going to the condenser (58 °C) up to the DSI temperature of 98 °C, the COP will decrease substantially to a COP of approximately 3 when compared to upgrading vapour from 73 °C (COP is approximately 10).

8.4.4 *Estimate the impact and iterate*

Once the heat pump retrofit bridge was identified and the electrical cost was determined, the fossil steam was replaced with upgraded vapour using the HSDT. The heat recovery potential and the resulting retrofit bridge was then removed from the diagram. Subsequently, a new HSDT of the remaining heat available for retrofit was formed, which also showed the degree of electrification for the new design. If the degree of electrification, the electrical cost and the number of retrofitted equipment required (each HSDT stream that is matched requires a new heat exchanger or heat pump to be retrofitted) was satisfactory when compared to the PE target set initially, no further iterations were required. The potential solution was then remodelled to determine its process viability. However, if the target was not satisfied then the resulting HSDT was used to determine a subsequent bridge match, furthering the electrification of the process. This is demonstrated in Figure 77.

For the example problem, the potential solution would consist of three new recovery exchangers and three MVRs in series to upgrade the vapour. The first two heat exchangers are integrated by splitting the excess vapour, one to heat the raw milk from 8 °C to 12 °C then another to heat the milk from 45 °C to 55 °C. The third exchanger purchased will use the COW water to heat the milk from 55 °C to 60 °C. Then, all 3,758 kW of high temperature excess vapour bleed from Effect 1 can then be upgraded to replace fossil steam in the DSI.

Please note that all HSDT and retrofit bridges for the three case studies are available in Appendix Four: All Heat Surplus-Deficit Tables for electrification retrofit.

As a summary, the steps taken to iterate through designs are as follows:

1. Produce the HSDT diagram.

2. Determine areas of under-utilised heat, for the example the cooling exchanger COND Lp has heat available at temperature (57 °C) that it can be exchanged with the sink of recovery exchanger Vap Bleed (45 °C–54 °C).
3. COW preheat also has a heat sink at a low enough temperature (8 °C–30 °C) to accept heat from COND Lp unlocking additional higher quality heat sources.
4. Utilise the unlocked heat to bridge to another lower temperature sink or upgrade the unlocked heat via heat pump technology. For the above example, COW preheats source is not a vapour (MVR cannot be used) therefore, a heat bridge is created to the Vap bleed sink, further unlocking high quality heat.
5. The unlocked 72 °C vapour is upgraded via MVR to replace the MP source used for DSI preheat.
6. Performance and costs are estimated.

8.4.5 *General retrofit solutions*

The results of applying the proposed retrofit methodology to the three case studies are shown in Table 25. Two distinctive retrofit pathways were common across all three cases with an example process flow diagram for each proposed solution shown in Figure 78, Figure 79. The first solution involves DSI replacement as shown in the example above Figure 77. The opportunity exists to take excess vapour, integrate recovery heat exchangers to unlock a higher temperature stream, then subsequently upgrade this vapour using MVRs for all three cases. In all cases, the DSI replacement with MVR units achieved a relatively low work consumption compared to a DSI replacement solution with an air sourced heat pump (ranging between 311 kW and 447 kW) whilst saving a large proportion of fossil steam (between 3,109 kW and 4,467 kW).

Table 25 Steam savings and electricity cost across the three case studies.

Design	Retrofit Method	Steam Savings Q'sav	Additional Electricity Pcom	Energy Savings	Emissions Savings
Two Effect MVR-TV	DSI Replacement	4467 kW	447 kW	38.2 %	44.6 %
Two Effect MVR-TV	TVR Replacement	4154 kW	126 kW	38.3 %	42.9 %
Three Effect MVR-TV	DSI Replacement	4175 kW	418 kW	37.0 %	43.2 %
Three Effect MVR-TV	TVR Replacement	3453 kW	115 kW	32.9 %	36.9 %
Four Effect TVR	DSI Replacement	3109 kW	311 kW	13.7 %	14.5 %
Four Effect TVR	TVR Replacement	10722 kW	2212 kW	41.6 %	47.3 %
DSI replacement	ΔT_{lift}	~ 23 °C	DSI replacement COP		~ 10
TVR replacement	ΔT_{lift}	~ 5 °C	TVR replacement COP		~ 50

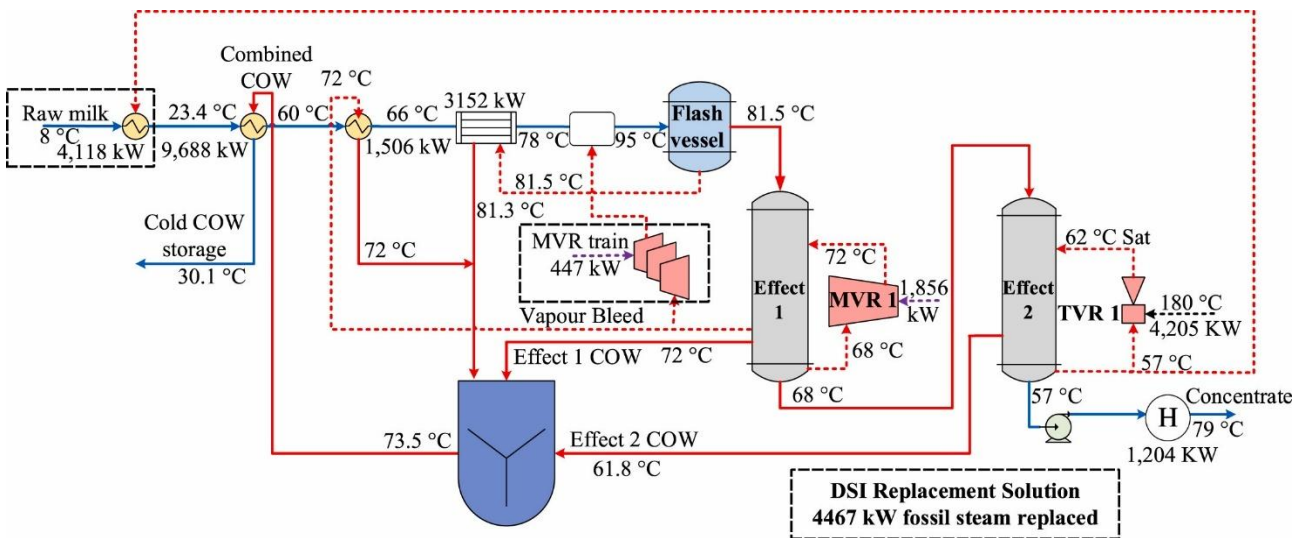


Figure 78 Example of DSI replacement for the two effect MVR-TV case.

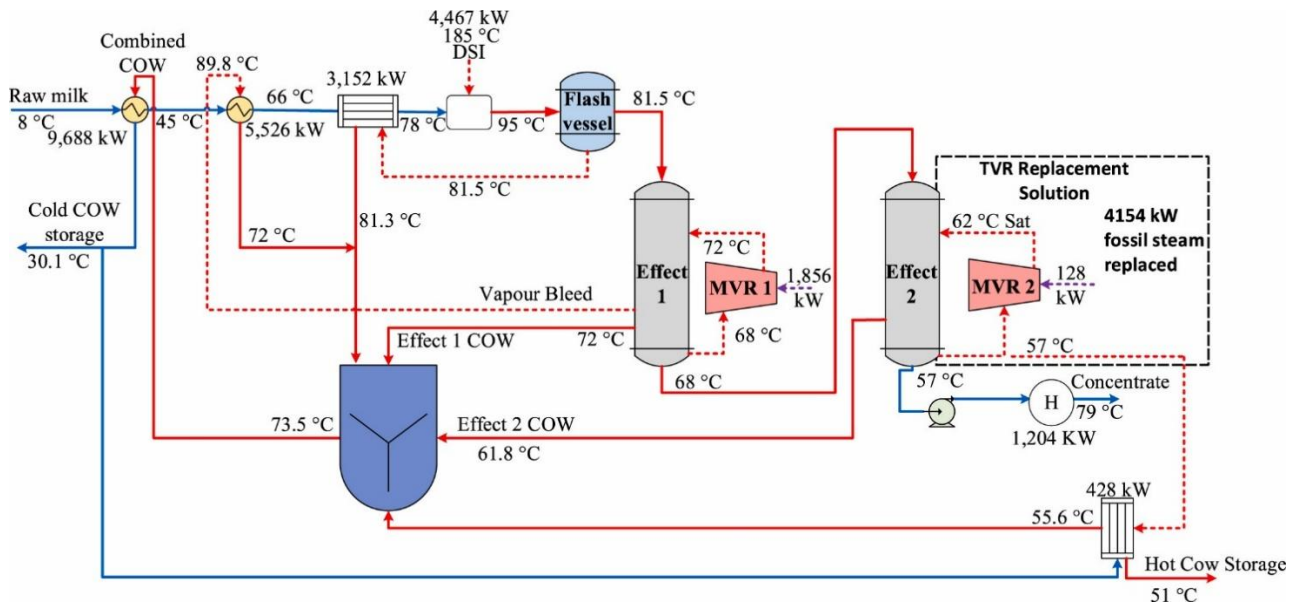


Figure 79 Example of TVR replacement solution for the two effect MVR-TVRC case.

The other viable retrofit strategy identified was TVR replacement. All cases presented involved heat surplus that was utilised in a TVR which relies on fossil steam (high-emitting process unit operation) for upgrading low-pressure vapour to satisfy the required heating for the effect. In the proposed retrofit strategy, the TVR can be replaced with an MVR, replacing the fossil steam with electricity. The efficiency of an MVR (COP of 50 assuming an isentropic efficiency of 86 %) is much larger than the efficiency of the TVR (equivalent to a COP of approximately 1.5). Therefore, less utility is required per unit of vapour to be upgraded. However, to achieve the required heating for the effect, a larger mass flowrate of low-pressure vapour is required to compensate for the lower utility and maintain the energy balance. Two out of three cases showed similar levels of steam savings for a significantly smaller electrical cost and number of modifications compared to the DSI replacement solution. The benefit was increased substantially for the four effect TVR case due to the reliance on fossil steam and TVR in older designs.

The retrofit strategies can also be used in conjunction with one another. The HSDT shows that the TVR replacement designs also have remaining heat surpluses available from the COND L_p stream ranging between 500 kW for the two effect MVR-TVRC, 1,000 kW for the three effect MVR-TVRC, and 8,000 kW for the four-effect design. These respective heat surpluses are sufficient to partially or fully replace the DSI, using the DSI replacement solution, effectively utilising both proposed viable retrofit strategies simultaneously.

8.4.6 *Comparison with conventional retrofit solutions*

Högnabba et al.(2024) presents one of the few recent dairy-specific heat pump retrofit studies, investigating the installation of a very high-temperature heat pump at a dairy factory in Maasdam, the Netherlands. Their system uses waste heat from an ice-water condenser, which is first upgraded to 85 °C using a hot-water heat pump and then supplied to a HighLift heat pump to generate 4 bar(g) steam at approximately 152 °C. The system was estimated to produce 625 kW of steam, with a top-cycle COP of 2.4 and an overall COP of 1.75 for the full lift from 35 °C waste heat to 152 °C steam.

This study demonstrates the feasibility and environmental benefit of replacing fossil-fuelled boiler steam with electrified steam generation. However, because the retrofit is framed as a steam replacement problem, the heat pump must perform a large overall temperature lift to produce steam at the required pressure. This constrains the achievable COP, even when staged heat pumping is used. In other words, the solution decarbonises the utility system, but it does not necessarily minimise the thermodynamic work required by the wider dairy process.

The method instead identifies where heat is required within the process, at what temperature level, and whether the demand can be met through lower-lift heat recovery or heat pump integration. This allows parts of the dairy process to be electrified using much smaller temperature lifts than would be required for full steam generation. In these cases, very high effective COPs can be achieved, including COPs greater than 10 for selected low-lift process heat duties. The advantage of this lower-lift targeting approach is that it avoids unnecessarily upgrading all recovered heat to steam temperature when only part of the process requires such high-grade heat.

8.4.7 *Economic assessment*

The economic assessment is based on the calculation of LCOH for different electricity-to-reference price ratios r_p . The economic parameters were assumed as shown Table 26. Capital costs for the MVRs are differentiated by mean values for different size classes and depending on the specific electric capacity of the MVRs. For the planning, installation and infrastructure measure necessary of the heat pump implementation, capital costs between 200 and 400 % of the investment costs were considered by the capital cost factor, $f_{C,Inst}$. Since the reference boiler is assumed as already existing and the retrofit of the heat pump is purely an efficiency measure, there are no capital costs for the

reference case. The investment costs are only used to calculate the operating costs of the reference system. All other parameters were assumed according to Schlosser et al (2023).

Table 26 Economic parameters for the LCOH calculation.

Parameter	MVR	Reference boiler
Investment costs, C_I	1,633 NZD/kW _{el} (for 100 – 1000 kW) 792 NZD/kW _{el} (for 1300 – 3000 kW)	237.65 NZD/kW _{th} (Schlosser, Zysk, et al., 2023)
Capital cost factor, $f_{C,Inst}$	2 to 4	
Reference energy carrier price, p_{ref}	$p_{el}=r_p \cdot p_{ref}$	$p_{ref}=55.39$ NZD/MWh
Full-load hours, t_{FL}	6000 h/a (three-shift operation)	
Interest factor, i	0.12	
Depreciation period, N	20 a	
Factor maintenance costs, $f_{M,j}$	2.5	3.0

Figure 80 shows that the DSI replacement variants differ only slightly in terms of nominal electrical compressor capacity and COP, which is why the LCOH are almost the same. The TVR replacement variants, by contrast, deviate from this. While the four-effect variant is the most expensive over the majority of all price ratios, this variant even exceeds the fossil reference variant for high capital costs ($f_{C,Inst} = 4$) and high price ratios. However, the other two cases (two and three effect MVR-TV R design) show cost advantages. In addition, they are less sensitive to the price ratio and capital costs due to their high efficiency and lower electrical compressor input.




Figure 80 LCOH of the three case studies for the TVR and DSI replacement design considering capital cost factors for planning, installation and infrastructure between $f_{C,Inst} = 2$ (solid lines) and $f_{C,Inst} = 4$ (dashed lines).

8.4.8 Other applications and limitations

Having a method that incorporates the processes heat cascade allows heat pump bridge analysis to be used in wider applications where self-recuperation is available such as other chemical industries. The method also assists in the integration of direct process electrification technologies, for example, a pulsed electric field into a potato blancher as the internal heat flows are known.

Currently, the method only considers local process zones or individual process retrofits, causing a limitation when needing to electrify a wider site with multiple zones. A combined method between bridge analysis and total site should be considered for future work.

8.5 Conclusions

As more pressure is placed on industries such as dairy to reduce emissions and the interest in process electrification grows, intuitive and easy to use retrofit methodologies and solutions will become a necessity. This study focused on refining modified Bridge Analysis into a Process Electrification Retrofit Methodology by including process unit operations and work consumption to reduce emissions in a cost-effective manner. To demonstrate the method, investigations were performed on three milk evaporator case studies, which led to the proposal of two general retrofit solutions. These solutions were applicable and cost effective across all cases with a TVR replacement being the simplest and most cost-effective solution overall, providing energy and emissions savings between 33.9 % - 47.6 % and 36.9 % - 47.3 % respectively as well as significant cost advantages across a wide range of electricity prices. The case studies showed that for effective PE, smart PI and correct technology selection is key. In all cases, the proposed replacement solutions were shown to be more cost effective than the reference system for electricity-to-reference price ratios up to 4.2. The case study also exhibited the requirement to consider all elements of a plant; including process unit operations in the retrofit study highlighted and quantified the most effective retrofit solution.

Overall, the Process Integration and Electrification Retrofit Method assisted in the discovery and quantification of retrofit solutions by allowing appropriate heat exchange and heat pump placement to be easily highlighted.

Chapter 9 Conclusions

This thesis has developed and demonstrated a comprehensive Process Integration and Electrification (PI&E) methodology supported by digital twins. Using the New Zealand dairy sector as a representative case study, the research established a systematic approach that supports both greenfield and retrofit contexts. The methodology integrates digital twins, advanced property modelling, exergy-based targeting, and techno-economic evaluation into a unified framework. Together, these elements provide a structured pathway to reduce fossil fuel reliance while ensuring designs remain technically feasible and economically viable.

The first contribution of the research was achieved through the application of the PI&E methodology to greenfield design. The thesis presented a fully electrified milk evaporation process as a case study,

combining exergy-based targeting, iterative operating point optimisation, and systematic heat pump integration. The study demonstrated how high-efficiency electrified systems can displace fossil-fuelled utilities, reduce exergy destruction, and improve overall energy performance. By grounding the analysis in realistic design constraints, including achievable compressor lift and minimum approach temperatures, the work showed that electrification can be both technically and economically competitive, particularly when integrated within a renewable-rich electricity system such as New Zealand's.

The addition of the NXLC is a supplementary contribution that addresses an interpretive limitation of conventional exergy Grand Composite Curves for systems operating across the dead-state temperature. By explicitly distinguishing exergy source and sink regions relative to the dead state, the NXLC provides a clearer basis for identifying work input, work recovery, and heat pump integration opportunities. This contribution further strengthens the exergy-based targeting component of the PI&E methodology.

The PI&E methodology was also extended to retrofit contexts, which represent the majority of industrial decarbonisation opportunities. A modified Bridge Analysis approach was introduced that adapts to existing heat exchanger networks. Multiple retrofit case studies were presented, ranging from subsystem reconfiguration to whole-plant electrification strategies. The results highlighted pathways where targeted interventions, such as replacement of thermal vapour recompression with mechanical alternatives, re-piping of condenser networks, and the integration of industrial heat pumps achieved substantial emissions reductions. These retrofit studies confirmed that incremental, cost-effective electrification is achievable even under the physical and economic constraints of existing plants.

The third contribution lies in digital twin development to accelerate PI&E. Chapter 4 detailed the creation of an open-source digital twin environment using the Ahuora platform, underpinned by the IDAES framework, and demonstrated its advantages for process integration and optimisation. In parallel, Design Digital Twins were constructed in commercial platforms such as Aspen HYSYS and DWSIM to benchmark performance and evaluate workflow interoperability. These efforts showed how Digital Twins enabled direct coupling between thermodynamic models, targeting methods, and optimisation tools, reducing manual data reconciliation and lowering barriers to the adoption of electrification workflows.

A further contribution is the synthesis of new methods that underpin the development of advanced thermophysical property models tailored to complex industrial fluids. The thesis demonstrated three complementary approaches for different types of fluids: activity coefficient models for multicomponent systems, such as food, Helmholtz-energy formulations for high-fidelity representation of pure refrigerants, and machine-learned surrogate models for computationally challenging mixtures such as humid air. These models contribute to closing a long-standing gap in the representation of non-ideal fluids, enabling accurate energy and exergy balances. Their integration into an equation-oriented framework ensured compatibility with an optimisation-driven digital twin environment, reducing uncertainty in electrification studies.

Collectively, these contributions establish a generalised and transferrable PI&E methodology. The framework advances beyond conventional process integration by explicitly addressing electrification challenges, incorporating accurate property models, and embedding digital tools for iterative design. By bridging the gap between academic methods and industrial practice, this work provides both the theoretical foundation and practical guidance required to accelerate the decarbonisation of low- to medium-temperature process industries.

9.1 Progress to Implementation

When the PI&E methodology was first developed, its application was carried out as a highly manual exercise using Microsoft Excel. This reflects the dominant practice among industrial engineers in New Zealand, where process integration studies are typically performed by exporting thermal stream data from process models or plant historians and reassembling them in spreadsheets for energy targeting. Such workflows, while accessible, are inherently labour-intensive and prone to error. Each iteration required data to be re-extracted, manually reconciled, and reformatted before integration analyses could be updated. This limited the scope for iterative design and hindered the evaluation of more advanced electrification technologies.

As the methodology matured, so too did the technical capability of the tools employed to support it. The workflow gradually shifted from being executed entirely in Excel to making use of established process simulators such as Aspen HYSYS and DWSIM. These environments provided a more rigorous basis for process modelling and allowed for direct integration of thermodynamic property packages. However, while they offered improvements over manual spreadsheet-based methods, these simulators lacked the flexibility and rapid feedback required for systematic PI&E studies. In particular, key methods described in this thesis, such as exergy-based pinch analysis were still absent,

meaning that external tools and manual reconciliation steps were still required to complete a full analysis.

The challenges encountered in the generalised implementation of the PI&E methodology is not uncommon as the literature is replete with examples of such manual iterations between different process analysis tools. To address these constraints, the author of this thesis, alongside Stephen Burroughs (another PhD student) and under the direction of my supervisor, co-led the development of a new free-to-use digital twin platform (The Ahuora Digital Twin Platform) that could serve as a common foundation for fully implementing advanced integration and electrification methods. This platform is designed not only to embed the PI&E methodology presented in this thesis but also to incorporate complementary methods developed within the wider Ahuora research programme, thereby creating a unified, extensible environment for process integration, optimisation, and decarbonisation studies.

Whilst still in beta development, it is crucial that methods such as PI&E are embedded within user-friendly tools if they are to deliver tangible impact in industry. To date, the implementation of PI&E within the Ahuora Platform is still a work in progress and part of the intended follow-on work from this thesis. More generally, the Ahuora Platform itself is a massive undertaking and a work in progress. The platform has been progressively tested and refined through engagement with practitioners. More than 40 engineers from across the New Zealand process industries have trialed the simulation and optimisation modules that form the core of the platform and provided structured feedback, using the milk powder plant case study as a demonstration model. These interactions have been invaluable in shaping the usability, functionality, and relevance of the platform, ensuring that it addresses real industrial needs rather than remaining a purely academic exercise. The architecture, design, and user feedback are issues that addressed by the thesis of Stephen Burroughs.

9.2 Limitations and Future Work

While this thesis has advanced the methodological and practical understanding of Process Integration and Electrification (PI&E), several limitations must be acknowledged. These limitations primarily reflect the inherent challenges of bridging detailed thermodynamic modelling with large-scale industrial application, as well as the boundaries of time, data, and scope within which the research was conducted. These limitations however can provide interesting opportunities for future work and therefore the thesis will conclude on discussion some of these potential options

Validation of results

A first limitation relates to the validation of results. The findings presented in this thesis are based primarily on simulation and modelling, supported by benchmarking against published data and representative case studies. Although sensitivity analyses were conducted to account for uncertainty, full-scale experimental validation at pilot or commercial scale was not feasible. Consequently, while the models provide strong theoretical and numerical evidence of feasibility, practical implementation may reveal additional challenges associated with equipment performance, plant operability, and site-specific integration.

As part of planned postdoctoral research, the author intends to deploy the PI&E methodology at operational industrial sites, working alongside plant engineers and equipment suppliers. This work will enable the translation of digital twin results into practice, providing direct evidence of how electrified technologies such as industrial heat pumps and vapour recompression systems perform under real conditions.

Greater Techno-economic Assessment

A second limitation lies in the techno-economic assessment of electrification technologies. While the research incorporated levelised cost of heat and comparative utility costs to evaluate the economic competitiveness of electrification pathways, the results remain sensitive to assumptions regarding electricity pricing, capital costs, and installation factors. In particular, the cost of industrial-scale heat pumps and vapour recompression systems can vary significantly across regions, suppliers, and project contexts. As a result, the economic results should be interpreted as indicative trends rather than absolute outcomes. Broader datasets on capital expenditure, maintenance costs, and learning-curve effects would strengthen confidence in the techno-economic findings.

To address this limitation, future work will focus on the development and integration of a dedicated techno-economic database into the Ahuora platform. This database will contain reference values for capital costs, operating costs, and installation factors relevant to the New Zealand industrial context, alongside regionally appropriate energy and electricity price datasets. By embedding these data directly within the digital twin environment, the PI&E methodology will be able to generate techno-economic evaluations for specific studies.

Large Variability of Food Properties

Another important limitation concerns the generalisability of fluid property models. Although this thesis developed advanced property packages for milk, refrigerants, and humid air, the scope of validation was restricted to available experimental datasets and published correlations. In the case of milk and other dairy fluids, measured thermophysical properties can vary by up to 10% due to compositional differences between samples. While the models developed in this research demonstrated close alignment with benchmark data, they cannot fully capture the variability present across different plants, regions, and seasons. As such, site-specific calibration or refinement would be required for industrial deployment.

Future research should therefore focus on expanding and strengthening the experimental data foundation underpinning these property models. This would involve systematic measurement campaigns across a range of dairy fluid compositions, including varying protein, fat, and mineral contents. Particular attention should be paid to dynamic behaviours such as boiling point elevation, heat capacity near phase transitions, and viscosity across concentration ranges. High-precision calorimetric and phase equilibrium experiments, conducted under controlled conditions, would provide critical reference data for refining and validating predictive models.

Beyond experimental work, there is also a need to explore hybrid modelling approaches that combine physics-based equations of state with machine learning methods capable of capturing non-ideal and composition-dependent behaviours. Physics-informed machine learning could be particularly valuable in interpolating across sparse datasets while retaining thermodynamic consistency. Such models could then be parameterised and calibrated using datasets collected from multiple dairy processing sites, ensuring broader applicability.

Total Site Considerations

A further limitation arises from the system boundaries adopted in the case studies. The focus of this thesis was primarily on the milk evaporation subsystem as a representative model of thermally intensive processes in the dairy sector. While this scope provided sufficient complexity to demonstrate the PI&E methodology, it does not encompass the full breadth of energy use in dairy powder plants, such as spray drying, refrigeration, and ancillary operations. Similarly, the total site integration of upstream and downstream processes was beyond the scope of this work. Extending the methodology to encompass whole-plant or multi-site systems would likely identify additional opportunities and constraints.

Accessibility

Finally, the accessibility of digital twin development represents an ongoing challenge. Although this research successfully demonstrated the use of open-source equation-oriented environments such as IDAES, the implementation required substantial expertise in thermodynamics, optimisation, and programming. For many industrial practitioners, the learning curve remains a significant barrier. While the workflow proposed in this thesis reduces complexity relative to traditional approaches, further efforts are needed to make digital twin creation accessible to a wider engineering audience.

Future work should therefore focus on developing user-friendly interfaces and workflow automation within platforms such as Ahuora. By embedding guided workflows, modular templates, and simplified property model integration, the digital twin environment can evolve from a tool for specialists into a practical resource for process engineers. Particular emphasis should be placed on automating routine tasks such as data import, flowsheet generation, and utility targeting, so that users can focus on evaluating decarbonisation options rather than configuring simulation environments.

Another avenue of research is the integration of educational and training resources directly into the platform. Interactive tutorials, example case studies, and preconfigured libraries of unit operations would help build user confidence and accelerate learning. Collaborations with industry and professional training organisations could ensure that these resources meet the needs of practitioners, bridging the gap between advanced research tools and day-to-day engineering practice.

Finally, efforts should be directed towards improving interoperability with existing industrial tools and plant data systems. By enabling seamless data exchange between digital twins, process historians, and common engineering software, users would be able to incorporate Digital Twins into existing workflows without requiring specialised programming expertise. This step would not only expand the usability of PI&E methods but also foster greater industrial adoption, ultimately supporting faster and more widespread deployment of electrification solutions.

In summary, the limitations identified in this thesis should not be seen as shortcomings but rather as catalysts for future research. Each presents an opportunity to advance the PI&E methodology beyond the scope of this work: validating designs in operational plants, embedding robust techno-economic data directly into digital twin platforms, expanding and refining thermophysical property models, extending analysis to full-site and multi-sector applications, and making digital twin tools more accessible to practitioners. Together, these directions outline an ambitious but necessary agenda for industrial decarbonisation research. While time constraints meant that only selected aspects could be

explored in this thesis, the next phase of work offers a clear and exciting pathway to strengthen, validate, and scale the methodology. These future steps will not only enhance technical confidence but also accelerate the transition from ad hoc simulation studies to widespread industrial adoption.

Chapter 10 References

- Abo-Sinna, M. A., Yousria Abo-Elnaga, Y., & Mousa, A. A. (2014). An interactive dynamic approach based on hybrid swarm optimization for solving multiobjective programming problem with fuzzy parameters. *Applied Mathematical Modelling*, *38*(7), 2000–2014. <https://doi.org/10.1016/j.apm.2013.10.013>
- Adamson, K.-M., Walmsley, T. G., Carson, J. K., Chen, Q., Schlosser, F., Kong, L., & Cleland, D. J. (2022). High-temperature and transcritical heat pump cycles and advancements: A review. *Renewable and Sustainable Energy Reviews*, *167*, 112798. <https://doi.org/10.1016/j.rser.2022.112798>
- Ahlqvist, V., Holmberg, P., & Tangerås, T. (2022). A survey comparing centralized and decentralized electricity markets. *Energy Strategy Reviews*, *40*, 100812. <https://doi.org/10.1016/j.esr.2022.100812>
- Ahmetović, E., Ibrić, N., Kravanja, Z., Grossmann, I. E., Maréchal, F., Čuček, L., & Kermani, M. (2018). Simultaneous optimisation and heat integration of evaporation systems including mechanical vapour recompression and background process. *Energy*, *158*, 1160–1191. <https://doi.org/10.1016/j.energy.2018.06.046>
- Amusat, O. O., Atia, A. A., Dudchenko, A. V., & Bartholomew, T. V. (2024). Modeling Framework for Cost Optimization of Process-Scale Desalination Systems with Mineral Scaling and Precipitation. *ACS ES&T Engineering*, *4*(5), 1028–1047. <https://doi.org/10.1021/acsestengg.3c00537>

- Arpagaus, C., Bless, F., Uhlmann, M., Schiffmann, J., & Bertsch, S. S. (2018). High temperature heat pumps: Market overview, state of the art, research status, refrigerants, and application potentials. *Energy*, *152*, 985–1010. <https://doi.org/10.1016/j.energy.2018.03.166>
- Aspelund, A., Berstad, D. O., & Gundersen, T. (2007). An Extended Pinch Analysis and Design procedure utilizing pressure based exergy for subambient cooling. *Applied Thermal Engineering, Selected Papers from the 9th Conference on Process Integration, Modelling and Optimisation for Energy Saving and Pollution Reduction – PRES2006*, *27*(16), 2633–2649. <https://doi.org/10.1016/j.applthermaleng.2007.04.017>
- AspenTech. (2025). *Aspen HYSYS | Leading Process Simulation Software for Oil & Gas | AspenTech*. <https://www.aspentech.com/en/products/engineering/aspens-hysys>
- Atuonwu, J., & Tassou, S. (2021). Decarbonisation of food manufacturing by the electrification of heat: A review of developments, technology options and future directions. *Trends in Food Science & Technology*, *107*, 168–182. <https://doi.org/10.1016/j.tifs.2020.10.011>
- Ba, L., Tangour, F., El Abbassi, I., & Absi, R. (2025). Analysis of Digital Twin Applications in Energy Efficiency: A Systematic Review. *Sustainability*, *17*(8), 3560. <https://doi.org/10.3390/su17083560>
- Becker, H. C. (Ed.). (2012). *Methodology and Thermo-Economic Optimization for Integration of Industrial Heat Pumps*. EPFL. <https://doi.org/10.5075/epfl-thesis-5341>
- Bell, I. H. (2024). Superancillary Equations for the Multiparameter Equations of State in REFPROP 10.0. *Journal of Physical and Chemical Reference Data*, *53*(1), 013102. <https://doi.org/10.1063/5.0191228>

- Bell, I. H., Lemmon, E. W., & Harvey, A. H. (2023). An Analysis of the Critical Region of Multiparameter Equations of State. *International Journal of Thermophysics*, 44(11), 158. <https://doi.org/10.1007/s10765-023-03261-8>
- Bell, I. H., Wronski, J., Quoilin, S., & Lemort, V. (2014). Pure and Pseudo-pure Fluid Thermophysical Property Evaluation and the Open-Source Thermophysical Property Library CoolProp. *Industrial & Engineering Chemistry Research*, 53(6), 2498–2508. <https://doi.org/10.1021/ie4033999>
- Bellemo, L., & Bergamini, R. (2022). Integration of high temperature CO₂ heat pumps and conventional heaters for spray dryers. *High Temperature Heat Pump Symposium*.
- Bon, J., Clemente, G., Vaquiro, H., & Mulet, A. (2010). Simulation and optimization of milk pasteurization processes using a general process simulator (ProSimPlus). *Computers & Chemical Engineering*, 34(3), 414–420. <https://doi.org/10.1016/j.compchemeng.2009.11.013>
- Bonhivers, J.-C., Alva-Argaez, A., Srinivasan, B., & Stuart, P. R. (2017). New analysis method to reduce the industrial energy requirements by heat-exchanger network retrofit: Part 2 – Stepwise and graphical approach. *Applied Thermal Engineering*, 119, 670–686. <https://doi.org/10.1016/j.applthermaleng.2015.05.085>
- Boukouvala, F., Misener, R., & Floudas, C. A. (2016). Global optimization advances in Mixed-Integer Nonlinear Programming, MINLP, and Constrained Derivative-Free Optimization, CDFO. *European Journal of Operational Research*, 252(3), 701–727. <https://doi.org/10.1016/j.ejor.2015.12.018>
- Bücker, D., & Wagner, W. (2006). Reference Equations of State for the Thermodynamic Properties of Fluid Phase n-Butane and Isobutane. *Journal of Physical and Chemical Reference Data*, 35(2), 929–1019. <https://doi.org/10.1063/1.1901687>

- Bühler, F., Zühlsdorf, B., Nguyen, T.-V., & Elmegaard, B. (2019). A comparative assessment of electrification strategies for industrial sites: Case of milk powder production. *Applied Energy*, 250, 1383–1401. <https://doi.org/10.1016/j.apenergy.2019.05.071>
- Burgard, A. P., Eason, J. P., Eslick, J. C., Ghouse, J. H., Lee, A., Biegler, L. T., & Miller, D. C. (2018). A Smooth, Square Flash Formulation for Equation-Oriented Flowsheet Optimization. In M. R. Eden, M. G. Ierapetritou, & G. P. Towler (Eds.), *13th International Symposium on Process Systems Engineering (PSE 2018)* (Vol. 44, pp. 871–876). Elsevier. <https://doi.org/10.1016/B978-0-444-64241-7.50140-3>
- Burroughs, S., Lincoln, B., Adeel, A., Severinsen, I., Lee, A., Amusat, O., Gunter, D., Nicholson, B., Apperley, M., Young, B., Siirola, J., & Walmsle, T. (2025). *New Directions and Software Tools Within the Process Systems Engineering Ecosystem*. 430–436. <https://doi.org/10.69997/sct.156838>
- Bynum, M. L., Hackebeil, G. A., Hart, W. E., Laird, C. D., Nicholson, B. L., Siirola, J. D., Watson, J.-P., & Woodruff, D. L. (2021). *Pyomo—Optimization Modeling in Python* (Vol. 67). Springer International Publishing. <https://doi.org/10.1007/978-3-030-68928-5>
- Chen, C., Lu, Y., & Banares-Alcantara, R. (2019). Direct and indirect electrification of chemical industry using methanol production as a case study. *Applied Energy*, 243, 71–90. <https://doi.org/10.1016/j.apenergy.2019.03.184>
- Choi, Y.-H. (1986). Effects of Temperature and Composition on the Thermal Conductivity and Thermal Diffusivity of Some Food Components. *Korean Journal of Food Science and Technology*, 18(5), 357–363.
- Dairy News Today. (2024). *New Zealand's Dairy Industry in 2023*. <https://dairynews.today/milkypedia/country/nz/>

- De Santis, R., Gironi, F., & Marrelli, L. (1976). Vapor-Liquid Equilibrium from a Hard-Sphere Equation of State. *Industrial & Engineering Chemistry Fundamentals*, 15(3), 183–189. <https://doi.org/10.1021/i160059a006>
- Deason, J., Wei, M., Leventis, G., Smith, S., & Schwartz, L. C. (2018). *Electrification of buildings and industry in the United States: Drivers, barriers, prospects, and policy approaches*. Lawrence Berkeley National Lab. (LBNL), Berkeley, CA (United States). <https://doi.org/10.2172/1430688>
- E. Örs, R. Schmidt, M. Mighani, & M. Shalaby. (2020). A Conceptual Framework for AI-based Operational Digital Twin in Chemical Process Engineering. *2020 IEEE International Conference on Engineering, Technology and Innovation (ICE/ITMC)*, 1–8. <https://doi.org/10.1109/ICE/ITMC49519.2020.9198575>
- EECA. (2019). *High-temperature Heat Pumps for low carbon process heating*. <https://genless.govt.nz/assets/Business-Resources/High-temperature-heat-pumps-for-low-carbon-process-heating.pdf>
- EECA. (2020a). *Fonterra | Coal boiler conversion*. EECA. <https://www.eeca.govt.nz/insights/case-studies-and-articles/fuel-switching-captures-economic-and-climate-benefits-for-fonterra/>
- EECA. (2020b). *McCain | Pulsed Electric Field*. EECA. <https://www.eeca.govt.nz/insights/case-studies-and-articles/mccain-pulsed-electric-field/>
- EECA. (2025). *Approved GIDI projects*. EECA. <https://www.eeca.govt.nz/co-funding-and-support/approved-gidi-projects/>

- Fenghour, A., Wakeham, W. A., Vesovic, V., Watson, J. T. R., Millat, J., & Vogel, E. (1995). The Viscosity of Ammonia. *Journal of Physical and Chemical Reference Data*, 24(5), 1649–1667. <https://doi.org/10.1063/1.555961>
- Ferguson, S. (2020). Apollo 13: The first digital twin. *Siemens® Htps://Blogs. Sw. Siemens. Com/Simcenter/Apollo-13-the-First-Digital-Twin.*
- Forrester, A., Sobester, A., & Keane, A. (2008). Constructing a Surrogate. In *Engineering Design via Surrogate Modelling* (pp. 33–76). John Wiley & Sons, Ltd. <https://doi.org/10.1002/9780470770801.ch2>
- Fredenslund, A., Jones, R. L., & Prausnitz, J. M. (1975). Group-contribution estimation of activity coefficients in nonideal liquid mixtures. *AIChE Journal*, 21(6), 1086–1099. <https://doi.org/10.1002/aic.690210607>
- Friso, D. (2022). Mathematical Modelling of the Entrainment Ratio of High Performance Supersonic Industrial Ejectors. *Processes*, 10(1), Article 1. <https://doi.org/10.3390/pr10010088>
- Future market Insights. (2024). *Milk Powder Market Size, Growth & Forecast 2025-2035.* <https://www.futuremarketinsights.com/reports/milk-powder-market>
- Gao, K., Wu, J., Bell, I., Harvey, A. H., & Lemmon, E. W. (2023). A Reference Equation of State with an Associating Term for Thermodynamic Properties of Ammonia. *NIST*, 52(1). <https://www.nist.gov/publications/reference-equation-state-associating-term-thermodynamic-properties-ammonia>
- Goldstein, D., Heyer, M., Jakobs, D., Schultz, E. S., & Biegler, L. T. (2022). Multilevel surrogate modeling of an amine scrubbing process for CO₂ capture. *AIChE Journal*, 68(6). <https://doi.org/10.1002/aic.17705>

- Grievies, M. (2014). Digital twin: Manufacturing excellence through virtual factory replication. *White Paper*, 1(2014), 1–7.
- Hamsani, M. N., Walmsley, T. G., Liew, P. Y., & Wan Alwi, S. R. (2018). Combined Pinch and exergy numerical analysis for low temperature heat exchanger network. *Energy*, 153, 100–112. <https://doi.org/10.1016/j.energy.2018.04.023>
- Herrmann, S., Kretschmar, H.-J., & Gatley, D. P. (2009). Thermodynamic Properties of Real Moist Air, Dry Air, Steam, Water, and Ice (RP-1485). *HVAC&R Research*, 15(5), 961–986. <https://doi.org/10.1080/10789669.2009.10390874>
- Högnabba, K., Tveit, T.-M., Zevenhoven, R., & Vittor, S. (2024). *De-carbonization of European Dairy industry – integration, performance estimates, and environmental impact assessment of a very high temperature heat pump.* 85–59. https://research.abo.fi/ws/portalfiles/portal/63948676/Hognabba_HTHP24_Extended_Abstr act.pdf
- Hu, J., Sari, O., Eicher, S., & Rija Rakotozanakajy, A. (2009). Determination of specific heat of milk at different fat content between 1 °C and 59 °C using micro DSC. *Journal of Food Engineering*, 90(3), 395–399. <https://doi.org/10.1016/j.jfoodeng.2008.07.009>
- Huber, M. L., Lemmon, E. W., Bell, I. H., & McLinden, M. O. (2022). The NIST REFPROP Database for Highly Accurate Properties of Industrially Important Fluids. *Industrial & Engineering Chemistry Research*, 61(42), 15449–15472. <https://doi.org/10.1021/acs.iecr.2c01427>
- Huber, M. L., Lemmon, E. W., Diky, V., Smith, B. L., & Bruno, T. J. (2008, August 8). *Chemically Authentic Surrogate Mixture Model for the Thermophysical Properties of a Coal-Derived Liquid Fuel* [Research-article]. American Chemical Society. (world). ACS Publications. <https://doi.org/10.1021/ef800314b>

Hyland, R., & Wexler, A. (1983). FORMULATIONS FOR THE THERMODYNAMIC PROPERTIES OF THE SATURATED PHASES OF H₂O FROM 173.15 TO 473.15 K. *Ashrae Transactions*. <https://www.semanticscholar.org/paper/FORMULATIONS-FOR-THE-THERMODYNAMIC-PROPERTIES-OF-OF-Hyland-Wexler/057d72ab7824cf5c6137fc84798256d02bbee748>

IAPWS. (2004). *Guideline on the Henry's Constant and Vapor-Liquid Distribution Constant for Gases in H₂O and D₂O at High Temperatures*. <https://iapws.org/documents/release/HenGuide>

IEA. (2018). *Industrial heat demand by temperature range, 2018 – Charts – Data & Statistics*. IEA. <https://www.iea.org/data-and-statistics/charts/industrial-heat-demand-by-temperature-range-2018>

IEA. (2022). *Key World Energy Statistics 2021 – Analysis*. IEA. <https://www.iea.org/reports/key-world-energy-statistics-2021>

IEA. (2024, March 1). *CO₂ Emissions in 2023 – Analysis*. IEA. <https://www.iea.org/reports/co2-emissions-in-2023>

IEA. (2025). *CO₂ Emissions – Global Energy Review 2025 – Analysis*. IEA. <https://www.iea.org/reports/global-energy-review-2025/co2-emissions>

Jacobsen, R. T., & Stewart, R. B. (1973). Thermodynamic Properties of Nitrogen Including Liquid and Vapor Phases from 63K to 2000K with Pressures to 10,000 Bar. *Journal of Physical and Chemical Reference Data*, 2(4). <https://doi.org/10.1063/1.3253132>

- Jeyamkondan, S., Jayas, D. S., & Holley, R. A. (1999). Pulsed Electric Field Processing of Foods: A Review. *Journal of Food Protection*, 62(9), 1088–1096. <https://doi.org/10.4315/0362-028X-62.9.1088>
- Kamel, D. A., Gadalla, M. A., Abdelaziz, O. Y., Labib, M. A., & Ashour, F. H. (2017). Temperature driving force (TDF) curves for heat exchanger network retrofit – A case study and implications. *Energy*, 123, 283–295. <https://doi.org/10.1016/j.energy.2017.02.013>
- Kang, L., & Liu, Y. (2015). Multi-objective optimization on a heat exchanger network retrofit with a heat pump and analysis of CO₂ emissions control. *Applied Energy*, 154, 696–708. <https://doi.org/10.1016/j.apenergy.2015.05.050>
- Kaviani, A., Aslani, A., Zahedi, R., Ahmadi, H., & Malekli, M. R. (2022). A new approach for energy optimization in dairy industry. *Cleaner Engineering and Technology*, 8, 100498. <https://doi.org/10.1016/j.clet.2022.100498>
- Klemeš, J. J., & Kravanja, Z. (2013). Forty years of Heat Integration: Pinch Analysis (PA) and Mathematical Programming (MP). *Current Opinion in Chemical Engineering, Biotechnology and Bioprocess Engineering / Process Systems Engineering*, 2(4), 461–474. <https://doi.org/10.1016/j.coche.2013.10.003>
- Kritzing, W., Karner, M., Traar, G., Henjes, J., & Sihn, W. (2018). Digital Twin in manufacturing: A categorical literature review and classification. *Ifac-PapersOnline*, 51(11), 1016–1022.
- Lal, N. S. (2020). *Automated retrofit of heat exchanger networks* [Thesis, The University of Waikato]. <https://researchcommons.waikato.ac.nz/handle/10289/13670>

- Lal, N., Walmsley, T., Walmsley, M., Atkins, M., & Neale, J. (2018). A Novel Heat Exchanger Network Bridge Retrofit Method using the Modified Energy Transfer Diagram. *Energy*, 155. <https://doi.org/10.1016/j.energy.2018.05.019>
- Lee, A., Eslick, J., Laird, C., Siirola, J., Zamarripa, M., Gunter, D., Shinn, J., Dowling, A., Bhattacharyya, D., Biegler, L., Burgard, A., & Miller, D. (2021). The IDAES process modeling framework and model library—Flexibility for process simulation and optimization. *Journal of Advanced Manufacturing and Processing*, 3. <https://doi.org/10.1002/amp2.10095>
- Leland, T. W., & Chappellear, P. S. (1968). THE CORRESPONDING STATES PRINCIPLE—A REVIEW OF CURRENT THEORY AND PRACTICE. *Industrial & Engineering Chemistry*, 60(7), 15–43. <https://doi.org/10.1021/ie50703a005>
- Lemmon, E. W., & Jacobsen, R. T. (2005). A New Functional Form and New Fitting Techniques for Equations of State with Application to Pentafluoroethane (HFC-125). *Journal of Physical and Chemical Reference Data*, 34(1), 69–108. <https://doi.org/10.1063/1.1797813>
- Lemmon, E. W., Jacobsen, R. T., Penoncello, S. G., & Friend, D. G. (2000). Thermodynamic Properties of Air and Mixtures of Nitrogen, Argon, and Oxygen From 60 to 2000 K at Pressures to 2000 Mpa. *NIST*, 29(No. 3). <https://www.nist.gov/publications/thermodynamic-properties-air-and-mixtures-nitrogen-argon-and-oxygen-60-2000-k-pressures>
- Lemmon, E. W., McLinden, M. O., & Wagner, W. (2009). Thermodynamic Properties of Propane. III. A Reference Equation of State for Temperatures from the Melting Line to 650 K and Pressures up to 1000 MPa. *Journal of Chemical & Engineering Data*, 54(12), 3141–3180. <https://doi.org/10.1021/jc900217v>

- Lemmon, E. W., & Span, R. (2015). Thermodynamic Properties of R-227ea, R-365mfc, R-115, and R-131I. *Journal of Chemical & Engineering Data*, 60(12), 3745–3758. <https://doi.org/10.1021/acs.jced.5b00684>
- Lewis, M. J. (1993). Physical Properties of Dairy Products. In R. K. Robinson (Ed.), *Modern Dairy Technology: Volume 2 Advances in Milk Products* (pp. 331–380). Springer US. https://doi.org/10.1007/978-1-4684-8172-3_5
- Lie, F. (2014). Review on Ejector Efficiencies in Various Ejector Systems. *Purdue University International Refrigeration and Air Conditioning Conference*. <https://docs.lib.purdue.edu/cgi/viewcontent.cgi?article=2532&context=iracc>
- Lincoln, B. J., Kong, L., Pineda, A. M., & Walmsley, T. G. (2022). Process integration and electrification for efficient milk evaporation systems. *Energy*, 258, 124885. <https://doi.org/10.1016/j.energy.2022.124885>
- Lincoln, B. J., Kong, L., Schlosser, F., & Walmsley, T. G. (2024). Process Integration and Electrification for retrofit: Case studies of milk evaporator systems. *Case Studies in Thermal Engineering*, 60, 104601. <https://doi.org/10.1016/j.csite.2024.104601>
- Linnhoff, B., & Flower, J. R. (1978). Synthesis of heat exchanger networks: I. Systematic generation of energy optimal networks. *AIChE Journal*, 24(4), 633–642. <https://doi.org/10.1002/aic.690240411>
- Linnhoff, B., & Hindmarsh, E. (1983). The pinch design method for heat exchanger networks. *Chemical Engineering Science*, 38(5), 745–763. [https://doi.org/10.1016/0009-2509\(83\)80185-7](https://doi.org/10.1016/0009-2509(83)80185-7)

- Ma, K., Sahinidis, N. V., Bindlish, R., Bury, S. J., Haghpanah, R., & Rajagopalan, S. (2022). Data-driven strategies for extractive distillation unit optimization. *Computers & Chemical Engineering*, *167*, 107970. <https://doi.org/10.1016/j.compchemeng.2022.107970>
- Madana Gopal, J. V., Morgan, R., De Sercey, G., & Vogiatzaki, K. (2023). Overview of Common Thermophysical Property Modelling Approaches for Cryogenic Fluid Simulations at Supercritical Conditions. *Energies*, *16*(2), 885. <https://doi.org/10.3390/en16020885>
- Madni, A. M., Madni, C. C., & Lucero, S. D. (2019). Leveraging digital twin technology in model-based systems engineering. *Systems*, *7*(1), 7.
- Madoumier, M., Azzaro-Pantel, C., Tanguy, G., & Gésan-Guiziou, G. (2015). Modelling the properties of liquid foods for use of process flowsheeting simulators: Application to milk concentration. *Journal of Food Engineering*, *164*, 70–89. <https://doi.org/10.1016/j.jfoodeng.2015.04.023>
- Majdisova, Z., & Skala, V. (2018). *A Radial Basis Function Approximation for Large Datasets*. <https://doi.org/10.48550/arXiv.1806.04243>
- March, L. (1998). *Introduction to Pinch Technology*. 63.
- Marmolejo-Correa, D., & Gundersen, T. (2013). New Graphical Representation of Exergy Applied to Low Temperature Process Design. *Industrial & Engineering Chemistry Research*, *52*(22), 7145–7156. <https://doi.org/10.1021/ie302541e>
- MBIE. (2016). *Unlocking our energy productivity and renewable potential—New Zealand Energy Efficiency and Conservation Strategy*. <https://www.mbie.govt.nz/dmsdocument/140-nzeecs-2017-2022-pdf>

- MBIE. (2019). *Decarbonising process heat* | Ministry of Business, Innovation & Employment.
<https://www.mbie.govt.nz/building-and-energy/energy-and-natural-resources/low-emissions-economy/decarbonising-process-heat/>
- MBIE. (2022a). *Dairy Manufacturing Fact Sheet*.
<https://www.mbie.govt.nz/assets/4805f01c9f/dairy-manufacturing-fact-sheet.pdf>
- MBIE. (2022b). *Electricity statistics* | Ministry of Business, Innovation & Employment.
<https://www.mbie.govt.nz/building-and-energy/energy-and-natural-resources/energy-statistics-and-modelling/energy-statistics/electricity-statistics/>
- McBride, K., & Sundmacher, K. (2019). Overview of Surrogate Modeling in Chemical Process Engineering. *Chemie Ingenieur Technik*, 91(3), 228–239.
<https://doi.org/10.1002/cite.201800091>
- Mckay, M. D., Beckman, R. J., & Conover, W. J. (2000). A Comparison of Three Methods for Selecting Values of Input Variables in the Analysis of Output From a Computer Code. *Technometrics*, 42(1), 55–61. <https://doi.org/10.1080/00401706.2000.10485979>
- Minim, L., Coimbra, J., & Minim, V. (2002). Influence of Temperature and Water and Fat Contents on the Thermophysical Properties of Milk. *Journal of Chemical & Engineering Data*, 47.
<https://doi.org/10.1021/je025546a>
- Misener, R., & Biegler, L. (2023). Formulating data-driven surrogate models for process optimization. *Computers & Chemical Engineering*, 179, 108411.
<https://doi.org/10.1016/j.compchemeng.2023.108411>

- Mohr, P. J., & Taylor, B. N. (2005). CODATA recommended values of the fundamental physical constants: 2002. *Reviews of Modern Physics*, 77(1), 1–107. <https://doi.org/10.1103/RevModPhys.77.1>
- Monogenidou, S., Assael, M., & Huber, M. (2018). Reference Correlation for the Thermal Conductivity of Ammonia from the Triple-Point Temperature to 680 K and Pressures up to 80 MPa. *Journal of Physical and Chemical Reference Data*, 47, 043101. <https://doi.org/10.1063/1.5053087>
- M.R.W. Walmsley, N.S. Lal, T.G. Walmsley, & M.J. Atkins. (2017). A modified energy transfer diagram for heat exchanger network retrofit bridge analysis. *Chemical Engineering Transactions*, 61, 907–912. <https://doi.org/10.3303/CET1761149>
- Munir, M. T., Yu, W., & Young, B. R. (2012). Recycle effect on the relative exergy array. *Chemical Engineering Research and Design*, 90(1), 110–118. <https://doi.org/10.1016/j.cherd.2011.06.015>
- Munir, M. T., Zhang, Y., Yu, W., Wilson, D. I., & Young, B. R. (2016). Virtual milk for modelling and simulation of dairy processes. *Journal of Dairy Science*, 99(5), 3380–3395. <https://doi.org/10.3168/jds.2015-10449>
- Nicholson, B., Siirola, J. D., Watson, J.-P., Zavala, V. M., & Biegler, L. T. (2018). pyomo.dae: A modeling and automatic discretization framework for optimization with differential and algebraic equations. *Mathematical Programming Computation*, 10(2), 187–223. <https://doi.org/10.1007/s12532-017-0127-0>
- Ning, C., & You, F. (2019). Optimization under uncertainty in the era of big data and deep learning: When machine learning meets mathematical programming. *Computers & Chemical Engineering*, 125, 434–448. <https://doi.org/10.1016/j.compchemeng.2019.03.034>

- Nordman, R., & Berntsson, T. (2009). Use of advanced composite curves for assessing cost-effective HEN retrofit I: Theory and concepts. *Applied Thermal Engineering*, 29(2), 275–281. <https://doi.org/10.1016/j.applthermaleng.2008.02.021>
- Padullés, R., Walmsley, T. G., Lincoln, B. J., Andersen, M. P., Jensen, J. K., & Elmegaard, B. (2024). Process integration and electrification through multiple heat pumps using a Lorenz efficiency approach. *Energy*, 311, 133348.
- Pal, P. K., Hens, A., Behera, N., & Lahiri, S. K. (2025). Digital twins: Transforming the chemical process industry—A review. *The Canadian Journal of Chemical Engineering*, 103(8), 3611–3636. <https://doi.org/10.1002/cjce.25611>
- Piller. (2022). *Compressors and blowers for open loop MVR cycle and steam heat pumps Piller Blowers and Compressors*. Annex 58.
- Piller. (2025). *Low flow MVR Blower—The PILLER VapoFan*. <https://www.piller.de/products-services/mvr-blower/vapofan/>
- Poulsen, L., & Zühlsdorf, B. (2022). *Task 1—Technologies*. Annex 58. <https://heatpumpingtechnologies.org/annex58/task1/>
- Qi, Q., Tao, F., Hu, T., Anwer, N., Liu, A., Wei, Y., Wang, L., & Nee, A. Y. C. (2021). Enabling technologies and tools for digital twin. *Journal of Manufacturing Systems, Digital Twin towards Smart Manufacturing and Industry 4.0*, 58, 3–21. <https://doi.org/10.1016/j.jmsy.2019.10.001>
- Rasheed, A., San, O., & Kvamsdal, T. (2020). Digital Twin: Values, Challenges and Enablers From a Modeling Perspective. *IEEE Access*, 8, 21980–22012. <https://doi.org/10.1109/ACCESS.2020.2970143>

- Renon, H., & Prausnitz, J. M. (1968). Local compositions in thermodynamic excess functions for liquid mixtures. *AIChE Journal*, *14*(1), 135–144. <https://doi.org/10.1002/aic.690140124>
- Rios, L. M., & Sahinidis, N. V. (2013). Derivative-free optimization: A review of algorithms and comparison of software implementations. *Journal of Global Optimization*, *56*(3), 1247–1293. <https://doi.org/10.1007/s10898-012-9951-y>
- RNZ. (2021). *NZ replacing coal boilers with wood pellets but some say it slows carbon neutral progress*. RNZ. <https://www.rnz.co.nz/news/national/453589/nz-replacing-coal-boilers-with-wood-pellets-but-some-say-it-slows-carbon-neutral-progress>
- Sargent, R. W. H. (1979). Flowsheeting. *Computers & Chemical Engineering*, *3*(1), 17–20. [https://doi.org/10.1016/0098-1354\(79\)80004-6](https://doi.org/10.1016/0098-1354(79)80004-6)
- Sarkar, J., Bhattacharyya, S., & Ram Gopal, M. (2005). Transcritical CO₂ heat pump systems: Exergy analysis including heat transfer and fluid flow effects. *Energy Conversion and Management*, *46*(13–14), 2053–2067. <https://doi.org/10.1016/j.enconman.2004.10.022>
- Schlosser, F., Arpagaus, C., & Walmsley, T. G. (2019). Heat Pump Integration by Pinch Analysis for Industrial Applications: A Review. *Chemical Engineering Transactions*, *76*, 7–12. <https://doi.org/10.3303/CET1976002>
- Schlosser, F., Jesper, M., Vogelsang, J., Walmsley, T. G., Arpagaus, C., & Hesselbach, J. (2020a). Large-scale heat pumps: Applications, performance, economic feasibility and industrial integration. *Renewable and Sustainable Energy Reviews*, *133*, 110219. <https://doi.org/10.1016/j.rser.2020.110219>
- Schlosser, F., Jesper, M., Vogelsang, J., Walmsley, T. G., Arpagaus, C., & Hesselbach, J. (2020b). Large-scale heat pumps: Applications, performance, economic feasibility and industrial

integration. *Renewable and Sustainable Energy Reviews*, 133, 110219.
<https://doi.org/10.1016/j.rser.2020.110219>

Schlosser, F., Walmsley, T. G., Lincoln, B., & Meschede, H. (2023). Retrofit of a galvanisation plant using advanced heat pump bridge analysis. *Chemical Engineering Transactions*, 103, 319–324.

Schlosser, F., Wiebe, H., Walmsley, T. G., Atkins, M. J., Walmsley, M. R. W., & Hesselbach, J. (2021). Heat Pump Bridge Analysis Using the Modified Energy Transfer Diagram. *Energies*, 14(1), Article 1. <https://doi.org/10.3390/en14010137>

Schlosser, F., Zysk, S., Walmsley, T. G., Kong, L., Zühlsdorf, B., & Meschede, H. (2023). Break-even of high-temperature heat pump integration for milk spray drying. *Energy Conversion and Management*, 291, 117304. <https://doi.org/10.1016/j.enconman.2023.117304>

Sleiti, A. K., Kapat, J. S., & Vesely, L. (2022). Digital twin in energy industry: Proposed robust digital twin for power plant and other complex capital-intensive large engineering systems. *Energy Reports*, 8, 3704–3726. <https://doi.org/10.1016/j.egy.2022.02.305>

Soares, R. M., Câmara, M. M., Feital, T., & Pinto, J. C. (2019). Digital Twin for Monitoring of Industrial Multi-Effect Evaporation. *Processes*, 7(8), 537. <https://doi.org/10.3390/pr7080537>

Son, H., Kim, M., & Kim, J.-K. (2022). Sustainable process integration of electrification technologies with industrial energy systems. *Energy*, 239, 122060.
<https://doi.org/10.1016/j.energy.2021.122060>

Span, R., & Wagner, W. (1996). A New Equation of State for Carbon Dioxide Covering the Fluid Region from the Triple-Point Temperature to 1100 K at Pressures up to 800 MPa. *Journal of Physical and Chemical Reference Data*, 25(6), 1509–1596. <https://doi.org/10.1063/1.555991>

- Sreepathi, B., & Rangaiah, G. P. (2014). *Review of Heat Exchanger Network Retrofitting Methodologies and Their Applications | Industrial & Engineering Chemistry Research*.
<https://pubs.acs.org/doi/10.1021/ie403075c>
- Stampfli, J. A., Atkins, M. J., Olsen, D. G., Walmsley, M. R. W., & Wellig, B. (2019). Practical heat pump and storage integration into non-continuous processes: A hybrid approach utilizing insight based and nonlinear programming techniques. *Energy*, *182*, 236–253.
<https://doi.org/10.1016/j.energy.2019.05.218>
- Takagi, T. (1996). Ultrasonic Speeds in Compressed Liquid and Vapor Pressures for Pentafluoroethane. *Journal of Chemical & Engineering Data*, *41*(6), 1325–1328.
<https://doi.org/10.1021/je960127x>
- Tao, F., Cheng, J., Qi, Q., Zhang, M., Zhang, H., & Sui, F. (2018). Digital twin-driven product design, manufacturing and service with big data. *The International Journal of Advanced Manufacturing Technology*, *94*(9), 3563–3576.
- Thol, M., & Lemmon, E. W. (2016). Equation of State for the Thermodynamic Properties of trans-1,3,3,3-Tetrafluoropropene [R-1234ze(E)]. *International Journal of Thermophysics*, *37*(3), 28.
<https://doi.org/10.1007/s10765-016-2040-6>
- Tillner-Roth, R., & Baehr, H. D. (1994). An International Standard Formulation for the Thermodynamic Properties of 1,1,1,2-Tetrafluoroethane (HFC-134a) for Temperatures from 170 K to 455 K and Pressures up to 70 MPa. *Journal of Physical and Chemical Reference Data*, *23*(5), 657–729. <https://doi.org/10.1063/1.555958>
- Tillner-Roth, R., & Yokozeki, A. (1997). An International Standard Equation of State for Difluoromethane (R-32) for Temperatures from the Triple Point at 136.34 K to 435 K and

Pressures up to 70 MPa. *Journal of Physical and Chemical Reference Data*, 26(6), 1273–1328.
<https://doi.org/10.1063/1.556002>

Timothy Gordon Walmsley. (2025). *openpinch/openpinch-toolkit: Open Pinch is an Excel spreadsheet that aids conducting a pinch analysis and total site analysis*.
<https://github.com/openpinch/openpinch-toolkit>

Tjoe, T. N., & Linnhoff, B. (1986). USING PINCH TECHNOLOGY FOR PROCESS RETROFIT. CHEMICAL ENGINEERING. *Chemical Engineering*, 93, 47–60.

Townsend, D. W., & Linnhoff, B. (1983). Heat and power networks in process design. Part I: Criteria for placement of heat engines and heat pumps in process networks. *AIChE Journal*, 29(5), 742–748. <https://doi.org/10.1002/aic.690290508>

Tsochatzidi, A., Arvanitidis, A. L., & Georgiadis, M. C. (2024). Model Based Optimization of Energy Consumption in Milk Evaporators. *Processes*, 12(1), 209.
<https://doi.org/10.3390/pr12010209>

Tyson, B., Drury, T., & Heck, D. (2010). Achieving low steam on cane with direct contact juice heaters. In R. C. Bruce (Ed.), *Proceedings of the 2010 Conference of the Australian Society of Sugar Cane Technologists held at Bundaberg, Queensland, Australia, 11-14 May 2010* (pp. 517–521). Australian Society of Sugar Cane Technologists. CABI Databases.

Wächter, A., & Biegler, L. T. (2006). On the Implementation of a Primal-Dual Interior Point Filter Line Search Algorithm for Large-Scale Nonlinear Programming. *Mathematical Programming*, 106(1), 25–57.

- Wagner, W., & Pruß, A. (2002). The IAPWS Formulation 1995 for the Thermodynamic Properties of Ordinary Water Substance for General and Scientific Use. *Journal of Physical and Chemical Reference Data*, 31(2), 387–535. <https://doi.org/10.1063/1.1461829>
- Wall, G. (1986). *Exergy—A useful concept*. <https://www.osti.gov/etdeweb/biblio/6956285>
- Wallin, E., Franck, P. Å., & Berntsson, T. (1990). Heat pumps in industrial processes—An optimization methodology. *Heat Recovery Systems and CHP*, 10(4), 437–446. [https://doi.org/10.1016/0890-4332\(90\)90092-X](https://doi.org/10.1016/0890-4332(90)90092-X)
- Walmsley, M. R. W., Walmsley, T. G., Atkins, M. J., & Neale, J. R. (2016). *Sustainable milk powder production using enhanced process integration and 100 % renewable energy*. 52, 559–564. <https://doi.org/10.3303/CET1652094>
- Walmsley, T. (2025). *Open Pinch* [Computer software].
- Walmsley, T. G., Atkins, M. J., Walmsley, M. R. W., & Neale, J. R. (2016). Appropriate placement of vapour recompression in ultra-low energy industrial milk evaporation systems using Pinch Analysis. *Energy, Green Strategy for Energy Generation and Saving towards Sustainable Development*, 116, 1269–1281. <https://doi.org/10.1016/j.energy.2016.04.026>
- Walmsley, T. G., Atkins, M. J., Walmsley, M. R. W., Philipp, M., & Peesel, R. H. (2018). Process and utility systems integration and optimisation for ultra-low energy milk powder production. *Energy*, 146, 67–81. <https://doi.org/10.1016/j.energy.2017.04.142>
- Walmsley, T., Lincoln, B., Padullés Solé, R., & Cleland, D. (2024). Advancing Industrial Process Electrification and Heat Pump Integration with New Exergy Pinch Analysis Targeting Techniques. *Energies*, 17, 2838. <https://doi.org/10.3390/en17122838>

- Wang Mengying, Deng Chun, Wang Yufei, Feng Xiao, & Lan Xingying. (2018). Process integration and selection of heat pumps in industrial processes. *Chemical Engineering Transactions*, 70, 1105–1110. <https://doi.org/10.3303/CET1870185>
- Wang, P., Yang, M., Peng, Y., Zhu, J., Ju, R., & Yin, Q. (2019). Sensor control in anti-submarine warfare—A digital twin and random finite sets based approach. *Entropy*, 21(8), 767.
- Wang, Z., & Hirai, S. (2011). Modeling and estimation of rheological properties of food products for manufacturing simulations. *Journal of Food Engineering*, 102(2), 136–144. <https://doi.org/10.1016/j.jfoodeng.2010.08.011>
- Wasser, H. P., Lombardi, P. A., Mattepu, S. Y., Richter, M., Komarnicki, P., & Pantaleo, A. M. (2023). Design of dairy systems as active Net-Zero Energy Factories. Technical and economic analysis of the German decarbonization process. *2023 8th IEEE Workshop on the Electronic Grid (eGRID)*, 1–6. <https://doi.org/10.1109/eGrid58358.2023.10380944>
- Westphalen, D. L., & Maciel, M. R. W. (2000). Pinch analysis of evaporation systems. *Brazilian Journal of Chemical Engineering*, 17, 525–538. <https://doi.org/10.1590/S0104-66322000000400017>
- Winchester, J. (2000). *Model based analysis of the operation and control of falling film evaporators: A thesis presented for the degree of Doctor of Philosophy in Technology and Engineering, Massey University, Palmerston North, New Zealand* [Massey University]. <http://hdl.handle.net/10179/2181>
- Wolfram Research Inc. (2024). *Mathematica, Version 14.2*. <https://www.wolfram.com/mathematica>
- Wong, T.-T., Luk, W.-S., & Heng, P.-A. (1997). Sampling with Hammersley and Halton Points. *Journal of Graphics Tools*, 2(2), 9–24. <https://doi.org/10.1080/10867651.1997.10487471>

- Wright, S. J. (1999). Continuous optimization (nonlinear and linear programming). *Foundations of Computer-Aided Process Design*, 7.
- Yang, M., Li, T., Feng, X., & Wang, Y. (2020). A simulation-based targeting method for heat pump placements in heat exchanger networks. *Energy*, 203, 117907. <https://doi.org/10.1016/j.energy.2020.117907>
- Y.Q. Lai, Z.A. Manan, & S.R. Wan Alwi. (2017). Heat exchanger network retrofit using individual stream temperature vs enthalpy plot. *Chemical Engineering Transactions*, 61, 1651–1656. <https://doi.org/10.3303/CET1761273>
- Yu, W., Patros, P., Young, B., Klinac, E., & Walmsley, T. G. (2022). Energy digital twin technology for industrial energy management: Classification, challenges and future. *Renewable and Sustainable Energy Reviews*, 161, 112407. <https://doi.org/10.1016/j.rser.2022.112407>
- Zhang, Y., Munir, M. T., Yu, W., & Young, B. R. (2014). Development of hypothetical components for milk process simulation using a commercial process simulator. *Journal of Food Engineering*, 121, 87–93. <https://doi.org/10.1016/j.jfoodeng.2013.07.033>

Chapter 11 Appendix

11.1 Appendix One: Code file snippets created for the Helmholtz formulation of Ammonia, n-butane and iso-butane.

```
1. """Predefined expression for Helmholtz EoS functions
2. """
3. #Extended from phi_ideal_type01.py by John Eslick
4. __author__ = "Ben Lincoln and Stephen Burroughs"
5.
6. import pyomo.environ as pyo
7.
```

```

8.
9. def phi_ideal_expressions_type04(model, parameters):
10.     """Type01 expression for the ideal part of dimensionless Helmholtz free energy
11.     Args:
12.         model (Block): Pyomo model
13.         parameters (dict): Main parameters dictionary
14.     Returns:
15.         dict: Expressions for ideal part of Helmholtz free energy
16.     """
17.     last_term = parameters["eos"]["last_term_ideal"]
18.     n0 = parameters["eos"]["n0"]
19.     Tc = parameters["basic"]["Tc"]
20.     g0 = parameters["eos"]["g0"]
21.     rng = range(4, last_term + 1)
22.     return {
23.         "phii": pyo.log(model.delta)
24.         + n0[1]
25.         + n0[2] * model.tau
26.         + (n0[3] - 1)
27.         * pyo.log(model.tau)
28.         + sum(n0[i] * pyo.log(1 - pyo.exp(-g0[i] * model.tau/Tc)) for i in rng),
29.         "phii_d": 1.0 / model.delta,
30.         "phii_dd": -1.0 / model.delta**2,
31.         "phii_t": n0[2]
32.         + (n0[3] - 1) / model.tau
33.         + sum(
34.             n0[i] * g0[i]
35.             / (Tc * (pyo.exp((model.tau * g0[i]) / Tc) - 1))
36.             for i in rng),
37.         "phii_tt": (1 - n0[3])/model.tau**2
38.         - sum(
39.             (n0[i] * g0[i]**2 * pyo.exp((model.tau * g0[i]) / Tc))
40.             / (Tc**2 * (pyo.exp((model.tau * g0[i]) / Tc) - 1)**2)
41.             for i in rng),
42.         "phii_dt": 0,
43.     }

```

```

1. """Predefined expression for Helmholtz EoS functions
2. """
3. #Extended from phi_residual_type02.py by John Eslick
4. __author__ = "Ben Lincoln and Stephen Burroughs"
5.
6. import pyomo.environ as pyo
7.
8.
9. def phi_residual_expressions_type05(model, parameters):
10.     """Type05 expression for the residual part of dimensionless Helmholtz free energy

```

```

11.     Args:
12.         model (Block): Pyomo model
13.         parameters (dict): Main parameters dictionary
14.     Returns:
15.         dict: Expressions for residual part of Helmholtz free energy
16.     """
17.     last_terms = parameters["eos"]["last_term_residual"]
18.     n = parameters["eos"]["n"]
19.     t = parameters["eos"]["t"]
20.     d = parameters["eos"]["d"]
21.     c = parameters["eos"]["c"]
22.     a = parameters["eos"]["a"]
23.     b = parameters["eos"]["b"]
24.     bi = parameters["eos"]["bi"]
25.     e = parameters["eos"]["e"]
26.     g = parameters["eos"]["g"]
27.     first_term = 1
28.     rng = []
29.     for last_term in last_terms:
30.         rng.append(range(first_term, last_term + 1))
31.         first_term = last_term + 1
32.     return {
33.         "phir": sum(n[i] * model.delta ** d[i] * model.tau ** t[i] for i in rng[0])
34.         + sum(
35.             n[i]
36.             * model.delta ** d[i]
37.             * model.tau ** t[i]
38.             * pyo.exp(-model.delta ** c[i])
39.             for i in rng[1]
40.         )
41.         + sum(
42.             n[i]
43.             * model.delta ** d[i]
44.             * model.tau ** t[i]
45.             * pyo.exp(
46.                 -a[i] * (model.delta - e[i]) ** 2 - b[i] * (model.tau - g[i]) ** 2
47.             )
48.             for i in rng[2]
49.         )
50.         + sum(
51.             n[i]
52.             * model.delta ** d[i]
53.             * model.tau ** t[i]
54.             * pyo.exp(
55.                 -a[i] * (model.delta - e[i]) ** 2 + 1/( b[i] * (model.tau - g[i]) ** 2+bi[i])
56.             )

```

```

57.         for i in rng[3]
58.     ),
59.     "phir_d": sum(n[i] * d[i] * model.delta ** (d[i] - 1) * model.tau ** t[i] for i in rng[0])
60. + sum(
61.     n[i]
62.     * pyo.exp(-model.delta ** c[i])
63.     * model.delta ** (d[i] - 1)
64.     * model.tau ** t[i]
65.     * (d[i] - c[i] * model.delta ** c[i])
66.     for i in rng[1]
67. )
68. + sum(
69.     n[i]
70.     * model.delta ** d[i]
71.     * model.tau ** t[i]
72.     * pyo.exp(
73.         -a[i] * (model.delta - e[i]) ** 2 - b[i] * (model.tau - g[i]) ** 2
74.     )
75.     * (d[i] / model.delta - 2 * a[i] * (model.delta - e[i]))
76.     for i in rng[2]
77. )
78. + sum(
79.     n[i]
80.     * model.delta ** d[i]
81.     * model.tau ** t[i]
82.     * pyo.exp(
83.         -a[i] * (model.delta - e[i]) ** 2 +1/( b[i] * (model.tau - g[i]) ** 2+bi[i])
84.     )
85.     * (d[i] / model.delta - 2 * a[i] * (model.delta-e[i]))
86.     for i in rng[3]
87. ),
88.     "phir_dd": sum(
89.         n[i] * d[i] * (d[i] - 1) * model.delta ** (d[i] - 2) * model.tau ** t[i]
90.         for i in rng[0]
91.     )
92. + sum(
93.     n[i]
94.     * pyo.exp(-model.delta ** c[i])
95.     * model.delta ** (d[i] - 2)
96.     * model.tau ** t[i]
97.     * (
98.         (d[i] - c[i] * model.delta ** c[i])
99.         * (d[i] - 1 - c[i] * model.delta ** c[i])
100.         - c[i] ** 2 * model.delta ** c[i]
101.     )
102.     for i in rng[1]

```

```

103.     )
104.     + sum(
105.         n[i]
106.         * model.tau ** t[i]
107.         * pyo.exp(
108.             -a[i] * (model.delta - e[i]) ** 2 - b[i] * (model.tau - g[i]) ** 2
109.         )
110.         * (
111.             -2.0 * a[i] * model.delta ** d[i]
112.             + 4 * a[i] ** 2 * model.delta ** d[i] * (model.delta - e[i]) ** 2
113.             - 4 * d[i] * a[i] * model.delta ** (d[i] - 1) * (model.delta - e[i])
114.             + d[i] * (d[i] - 1) * model.delta ** (d[i] - 2)
115.         )
116.         for i in rng[2]
117.     )
118.     + sum(
119.         n[i]
120.         * model.delta ** (d[i]-2)
121.         * model.tau ** t[i]
122.         * (
123.             d[i] **2 + d[i]
124.             * (-4 * model.delta * a[i] * (model.delta - e[i]) -1)
125.             + 2 * model.delta ** 2 * a[i]
126.             * (2 * a[i] * (model.delta - e[i]) ** 2 - 1)
127.         )
128.         * pyo.exp(
129.             -a[i] * (model.delta - e[i]) ** 2 + 1 / ( b[i] * (model.tau - g[i]) ** 2 + bi[i])
130.         )
131.         for i in rng[3]
132.     ),
133.     "phir_t": sum(
134.         n[i] * t[i] * model.delta ** d[i] * model.tau ** (t[i] - 1) for i in rng[0]
135.     )
136.     + sum(
137.         n[i]
138.         * t[i]
139.         * model.delta ** d[i]
140.         * model.tau ** (t[i] - 1)
141.         * pyo.exp(-model.delta ** c[i])
142.         for i in rng[1]
143.     )
144.     + sum(
145.         n[i]
146.         * model.delta ** d[i]
147.         * model.tau ** t[i]
148.         * pyo.exp(

```

```

149.         -a[i] * (model.delta - e[i]) ** 2 - b[i] * (model.tau - g[i]) ** 2
150.     )
151.     * (t[i] / model.tau - 2 * b[i] * (model.tau - g[i]))
152.     for i in rng[2]
153. )
154. + sum(
155.     (n[i]
156.     * model.delta ** d[i]
157.     * model.tau ** (t[i] - 1)
158.     * (bi[i] ** 2 * t[i]
159.     + 2 * b[i] * bi[i] * t[i] * (model.tau - g[i]) ** 2
160.     - b[i] * (model.tau - g[i])
161.     * (b[i] * t[i] * (model.tau - g[i]) ** 3 + 2 * model.tau))
162.     * pyo.exp(
163.     -a[i] * (model.delta - e[i]) ** 2 + 1/( b[i] * (model.tau - g[i]) ** 2+bi[i])
164.     ))
165.     / (b[i] * (model.tau - g[i]) ** 2 + bi[i]) ** 2
166.     for i in rng[3]
167. ),
168. "phir_tt": sum(
169.     n[i] * t[i] * (t[i] - 1) * model.delta ** d[i] * model.tau ** (t[i] - 2)
170.     for i in rng[0]
171. )
172. + sum(
173.     n[i]
174.     * t[i]
175.     * (t[i] - 1)
176.     * model.delta ** d[i]
177.     * model.tau ** (t[i] - 2)
178.     * pyo.exp(-model.delta ** c[i])
179.     for i in rng[1]
180. )
181. + sum(
182.     n[i]
183.     * model.delta ** d[i]
184.     * model.tau ** t[i]
185.     * pyo.exp(
186.     -a[i] * (model.delta - e[i]) ** 2 - b[i] * (model.tau - g[i]) ** 2
187.     )
188.     * (
189.     (t[i] / model.tau - 2 * b[i] * (model.tau - g[i])) ** 2
190.     - t[i] / model.tau**2
191.     - 2 * b[i]
192.     )
193.     for i in rng[2]
194. )

```

```

195.     + sum(
196.         n[i]
197.         * model.delta ** d[i]
198.         * model.tau ** t[i]
199.         * ((4 * b[i] ** 2 * (model.tau - g[i]) ** 2
200.            * (2 * b[i] * (model.tau - g[i]) ** 2 + 2 * bi[i] + 1))
201.            / (b[i] * (model.tau - g[i]) ** 2 + bi[i]) ** 4
202.            - (2 * b[i] * (-2 * g[i] * t[i] + 2 * t[i] * model.tau + model.tau))
203.            / (model.tau * (b[i] * (model.tau - g[i]) ** 2 + bi[i]) ** 2)
204.            + ((t[i] - 1) * t[i])
205.            / model.tau ** 2
206.         )
207.         * pyo.exp(
208.             -a[i] * (model.delta - e[i]) ** 2 + 1 / (b[i] * (model.tau - g[i]) ** 2 + bi[i])
209.         )
210.     for i in rng[3]
211. ),
212. "phir_dt": sum(
213.     n[i] * t[i] * d[i] * model.delta ** (d[i] - 1) * model.tau ** (t[i] - 1)
214.     for i in rng[0]
215. )
216. + sum(
217.     n[i]
218.     * t[i]
219.     * model.delta ** (d[i] - 1)
220.     * model.tau ** (t[i] - 1)
221.     * (d[i] - c[i] * model.delta ** c[i])
222.     * pyo.exp(-model.delta ** c[i])
223.     for i in rng[1]
224. )
225. + sum(
226.     n[i]
227.     * model.delta ** d[i]
228.     * model.tau ** t[i]
229.     * pyo.exp(
230.         -a[i] * (model.delta - e[i]) ** 2 - b[i] * (model.tau - g[i]) ** 2
231.     )
232.     * (d[i] / model.delta - 2.0 * a[i] * (model.delta - e[i]))
233.     * (t[i] / model.tau - 2.0 * b[i] * (model.tau - g[i]))
234.     for i in rng[2]
235. )
236. + sum(
237.     n[i]
238.     * model.delta ** (d[i] - 1)
239.     * model.tau ** (t[i] - 1)
240.     * 1 / (b[i] * (model.tau - g[i]) ** 2 + bi[i]) ** 2

```

```

241.         * (d[i] - 2 * model.delta * a[i] * (model.delta - e[i]))
242.         * (bi[i] ** 2 * t[i]
243.           + 2 * b[i] * bi[i] * t[i] * (model.tau - g[i]) ** 2
244.           - b[i] * (model.tau - g[i])
245.           * (b[i] * t[i] * (model.tau - g[i]) ** 3 + 2 * model.tau))
246.         * pyo.exp(
247.           -a[i] * (model.delta - e[i]) ** 2 +1/( b[i] * (model.tau - g[i]) ** 2+bi[i])
248.         )
249.         for i in rng[3]
250.     ),
251. }
252.

```

```

1. {
2.     "comp": "nh3",
3.     "basic": {
4.         "R": 0.48820939114,
5.         "MW": 17.03052,
6.         "T_star": 405.56,
7.         "rho_star": 233.25,
8.         "Tc": 405.56,
9.         "rhoc": 233.25,
10.        "Pc": 11363.4,
11.        "Tt": 195.49,
12.        "Pt": 6.055813,
13.        "rhot_l": 733.8486491,
14.        "rhot_v": 0.063731569,
15.        "P_min": 6.055813,
16.        "P_max": 1000000.0,
17.        "rho_max": 904.8315276,
18.        "T_min": 195.49,
19.        "T_max": 725.0
20.    },
21.    "eos": {
22.        "reference": [
23.            "SKehui Gao (高克慧),Jiangtao Wu (吴江涛),Ian H. Bell, Allan H. Harvey, and Eric W.
Lemmon (2023)",
24.            "    A Reference Equation of State with an Associating Term for Thermodynamic
Properties of Ammonia"
25.        ],
26.        "c": {
27.            "6": 2,
28.            "7": 2,
29.            "8": 1
30.        },
31.        "d": {
32.            "1": 4,

```

```
33.         "2": 1,
34.         "3": 1,
35.         "4": 2,
36.         "5": 3,
37.         "6": 3,
38.         "7": 2,
39.         "8": 3,
40.         "9": 1,
41.         "10": 1,
42.         "11": 1,
43.         "12": 2,
44.         "13": 2,
45.         "14": 1,
46.         "15": 3,
47.         "16": 3,
48.         "17": 1,
49.         "18": 1,
50.         "19": 1,
51.         "20": 1
52.
53.     },
54.     "t": {
55.         "1": 1.0,
56.         "2": 0.382,
57.         "3": 1.00,
58.         "4": 1.0,
59.         "5": 0.677,
60.         "6": 2.915,
61.         "7": 3.51,
62.         "8": 1.063,
63.         "9": 0.655,
64.         "10": 1.3,
65.         "11": 3.1,
66.         "12": 1.4395,
67.         "13": 1.623,
68.         "14": 0.643,
69.         "15": 1.13,
70.         "16": 4.5,
71.         "17": 1.0,
72.         "18": 4.0,
73.         "19": 4.3315,
74.         "20": 4.015
75.     },
76.     "n": {
77.         "1": 0.006132232,
78.         "2": 1.7395866,
```

```
79.         "3": -2.2261792,
80.         "4": -0.30127553,
81.         "5": 0.08967023,
82.         "6": -0.076387037,
83.         "7": -0.84063963,
84.         "8": -0.27026327,
85.         "9": 6.212578,
86.         "10": -5.7844357,
87.         "11": 2.4817542,
88.         "12": -2.3739168,
89.         "13": 0.01493697,
90.         "14": -3.7749264,
91.         "15": 0.0006254348,
92.         "16": -1.7359e-05,
93.         "17": -0.13462033,
94.         "18": 0.07749072839,
95.         "19": -1.6909858,
96.         "20": 0.93739074
97.     },
98.     "a": {
99.         "9": 0.42776,
100.        "10": 0.6424,
101.        "11": 0.8175,
102.        "12": 0.7995,
103.        "13": 0.91,
104.        "14": 0.3574,
105.        "15": 1.21,
106.        "16": 4.14,
107.        "17": 22.56,
108.        "18": 22.68,
109.        "19": 2.8452,
110.        "20": 2.8342
111.    },
112.     "b": {
113.         "9": 1.708,
114.         "10": 1.4865,
115.         "11": 2.0915,
116.         "12": 2.43,
117.         "13": 0.488,
118.         "14": 1.1,
119.         "15": 0.85,
120.         "16": 1.14,
121.         "17": 945.64,
122.         "18": 993.85,
123.         "19": 0.3696,
124.         "20": 0.2962
```

```

125.     },
126.     "bi":{
127.         "19": 1.244,
128.         "20": 0.6826
129.     },
130.     "g": {
131.         "9": 1.036,
132.         "10": 1.2777,
133.         "11": 1.083,
134.         "12": 1.2906,
135.         "13": 0.928,
136.         "14": 0.934,
137.         "15": 0.919,
138.         "16": 1.852,
139.         "17": 1.05897,
140.         "18": 1.05277,
141.         "19": 1.108,
142.         "20": 1.313
143.     },
144.     "e": {
145.         "9": -0.0726,
146.         "10": -0.1274,
147.         "11": 0.7527,
148.         "12": 0.57,
149.         "13": 2.2,
150.         "14": -0.243,
151.         "15": 2.96,
152.         "16": 3.02,
153.         "17": 0.9574,
154.         "18": 0.9576,
155.         "19": 0.4478,
156.         "20": 0.44689
157.     },
158.     "reference_state_offset": [
159.
160.     ],
161.     "n0": {
162.         "1": -6.59406093943886,
163.         "2": 5.60101151987913,
164.         "3": 4.0,
165.         "4": 2.224,
166.         "5": 3.148,
167.         "6": 0.9579
168.     },
169.     "g0": {
170.         "4":1646,

```

```

171.         "5":3965,
172.         "6":7231
173.     },
174.     "last_term_ideal": 6,
175.     "last_term_residual": [5, 8, 18, 20],
176.     "phi_ideal_type": 4,
177.     "phi_residual_type": 5
178. },
179. "aux": {
180.     "reference": [
181.         "Gao, Kehui, et al. A reference equation of state with an associating term for the
thermodynamic properties of ammonia.",
182.         "    Journal of Physical and Chemical Reference Data 52.1 (2023).",
183.     ],
184.     "delta_l_sat_approx": {
185.         "c": 1,
186.         "n": {
187.             "1": 0.051236,
188.             "2": 3.7925,
189.             "3": -3.5929,
190.             "4": 4.6409,
191.             "5": -1.9893,
192.             "6": 1.5978
193.         },
194.         "t": {
195.             "1": 0.07,
196.             "2": 0.46,
197.             "3": 0.77,
198.             "4": 1.05,
199.             "5": 1.25,
200.             "6": 8.0
201.         },
202.         "type": 1
203.     },
204.     "delta_v_sat_approx": {
205.         "c": 1,
206.         "n": {
207.             "1": -0.089966,
208.             "2": -3.8722,
209.             "3": -8.1183,
210.             "4": -25.293,
211.             "5": -54.279,
212.             "6": -400.83
213.         },
214.         "t": {
215.             "1": 0.112,

```

```

216.         "2": 0.473,
217.         "3": 1.5,
218.         "4": 3.875,
219.         "5": 8.0,
220.         "6": 20.0
221.     },
222.     "type": 2
223. }
224. },
225. "transport": {
226.     "thermal_conductivity": {
227.         "reference": [
228.             "Monogenidou, S. A., Marc J. Assael, and Marcia L. Huber.",
229.             "Reference correlation for the thermal conductivity of ammonia ",
230.             "from the triple-point temperature to 680 K and pressures up to 80 MPa. ",
231.             "Journal of Physical and Chemical Reference Data 47.4 (2018).",
232.         ]
233.     },
234.     "viscosity": {
235.         "reference": [
236.             "Fenghour, A., et al. The viscosity of ammonia.",
237.             "Journal of Physical and Chemical Reference Data 24.5 (1995): 1649-1667.",
238.         ]
239.     },
240.     "surface_tension": {
241.         "reference": [
242.             "Mulero, A., I. Cachadina, Parra, M., Recommended Correlations for the Surface ",
243.             "Tension of Common Fluids, J. Phys. Chem. Ref. Data 41, 043105, (2012).",
244.         ],
245.         "Tc": 405.4,
246.         "s": {
247.             "0": 102.8,
248.             "1": -94.53
249.         },
250.         "n": {
251.             "0": 1.211,
252.             "1": 5.585
253.         },
254.         "type": 1
255.     }
256. }
257. }
258.

```

1. ""Generate parameter and expression files for NH3

```

2.
5.
6. import os
7. import pyomo.environ as pyo
8. from pyomo.common.fileutils import this_file_dir
9. from idaes.models.properties.general_helmholtz.helmholtz_parameters import (
10.     WriteParameters,
11. )
12.
13.
14. def thermal_conductivity_rule(m):
15.     # """Thermal conductivity rule
16.
17.     # Monogenidou, S. A., Marc J. Assael, and Marcia L. Huber.
18.     # "Reference correlation for the thermal conductivity of ammonia
19.     # from the triple-point temperature to 680 K and pressures up to 80 MPa."
20.     # Journal of Physical and Chemical Reference Data 47.4 (2018).
21.     # """
22.     b = {
23.         1: 0.07152,
24.         2: 130.228,
25.         3: 9569.817
26.     }
27.     u = {
28.         1: 1646,
29.         2: 3965,
30.         3: 7231
31.     }
32.     v = {
33.         1: 2.224,
34.         2: 3.148,
35.         3: 0.9579
36.     }
37.     B = {
38.         1: {
39.             1: 0.103432e0,
40.             2: -0.112597e0,
41.             3: 0.233301e0,
42.             4: -0.112536e0,
43.             5: 0.141129e-1
44.         },
45.         2: {
46.             1: -0.283976e-1,
47.             2: 0.482520e-1,
48.             3: -0.644124e-1,
49.             4: 0.529376e-2,

```

```

50.         5: 0.891203e-2
51.     }
52. }
53. T = m.T_star / m.tau
54. # Ts = T / 251.196
55. rho = m.rho_star * m.delta
56. G = sum(bval / T**i for i, bval in b.items())
57. cint_over_k = 4.0
58. + sum(v[i] * (u[i] / T) ** 2
59.     * pyo.exp(u[i]/T)
60.     / (pyo.exp(u[i]/T)-1) ** 2
61.     for i in range(1,len(u.items()+1)
62. )
63. return (
64.     0.1351767 * pyo.sqrt(T) * cint_over_k / G
65.     +sum((B[1][i] + B[2][i]*T)*rho ** i for i in range(1,6))
66. )
67.
68.
69. def viscosity_rule(m):
70.     # """Viscosity rule
71.
72.     # Fenghour, A., et al. "The viscosity of ammonia."
73.     # Journal of Physical and Chemical Reference Data 24.5 (1995): 1649-1667.
74.     # """
75.     a = {
76.         0: 4.99318220,
77.         1: -0.61122364,
78.         2: 0,
79.         3: 0.18535124,
80.         4: -0.11160946,
81.     }
82.     c = {
83.         0: -0.17999496e1,
84.         1: 0.46692621e2,
85.         2: -0.53460794e3,
86.         3: 0.33604074e4,
87.         4: -0.13019164e5,
88.         5: 0.33414230e5,
89.         6: -0.58711743e5,
90.         7: 0.71426686e5,
91.         8: -0.59834012e5,
92.         9: 0.33652741e5,
93.         10: -0.12027350e5,
94.         11: 0.24348205e4,
95.         12: -0.20807957e3

```

```

96.     }
97.     d = {
98.         2: {
99.             0: 0,
100.            1: 0,
101.            2: 2.19664285e-1,
102.            3: 0,
103.            4: -0.83651107e-1
104.        },
105.        3: {
106.            0: 0.17366936e-2,
107.            1: -0.64250359e-2,
108.            2: 0,
109.            3: 0,
110.            4: 0
111.        },
112.        4: {
113.            0: 0,
114.            1: 0,
115.            2: 1.67668649e-4,
116.            3: -1.49710093e-4,
117.            4: 0.77012274e-4
118.        }
119.    }
120.
121.    T = m.T_star / m.tau
122.    rho = m.delta * m.rho_star
123.    Ts = T / 386
124.    return (
125.        0.021357
126.        * pyo.sqrt(17.03026*T)
127.        / (0.2957 ** 2 * pyo.exp(sum(aval * pyo.log(Ts) ** i for i, aval in a.items()))))
128.        +
129.        0.021357
130.        * pyo.sqrt(17.03026*T)
131.        / (0.2957 ** 2 * pyo.exp(sum(aval * pyo.log(Ts) ** i for i, aval in a.items()))))
132.        * sum((c[i] * pyo.sqrt(Ts)**-i)
133.              for i in range(0, len(c.items())))
134.        / (0.6022137 * 0.2957 **3)
135.        * rho
136.        + sum(sum(d[i][j]/Ts**j for j in range(0, len(d[i].items())))
137.              * rho ** i for i in range(2,5))
138.    )
139.
140.
141. def main(dry_run=False):

```

```

142.     """Generate parameter and expression files.
143.     Args:
144.         dry_run (bool): If dry run don't generate files
145.     Returns:
146.         None
147.     """
148.     main_param_file = os.path.join(this_file_dir(), "nh3.json")
149.     we = WriteParameters(parameters=main_param_file)
150.     we.add(
151.         {
152.             "viscosity": viscosity_rule,
153.             "thermal_conductivity": thermal_conductivity_rule,
154.         }
155.     )
156.     we.write(dry_run=dry_run)
157.     return we
158.
159.
160. if __name__ == "__main__":
161.     main()
162.

```

```

1. {
2.     "comp": "butane",
3.     "basic": {
4.         "R": 0.14305157,
5.         "MW": 58.12220,
6.         "T_star": 425.125,
7.         "rho_star": 227.8,
8.         "Tc": 425.125,
9.         "rhoc": 228.0,
10.        "Pc": 3796.0,
11.        "Tt": 134.895,
12.        "Pt": 0.000653,
13.        "rhot_l": 734.96,
14.        "rhot_v": 0.0000345,
15.        "P_min": 0.000653,
16.        "P_max": 69e6,
17.        "rho_max": 750,
18.        "T_min": 134.895,
19.        "T_max": 575
20.    },
21.    "eos": {
22.        "reference": [
23.            "Reference Equations of State for the Thermodynamic",
24.            "Properties of Fluid Phase -Butane and Isobutane",
25.            "D. Bucker; W. Wagner"

```

```
26.     ],
27.     "c": {
28.         "8": 1,
29.         "9": 1,
30.         "10": 1,
31.         "11": 1,
32.         "12": 1,
33.         "13": 1,
34.         "14": 2,
35.         "15": 2,
36.         "16": 2,
37.         "17": 2,
38.         "18": 2,
39.         "19": 2,
40.         "20": 2,
41.         "21": 2,
42.         "22": 3,
43.         "23": 3
44.     },
45.     "d": {
46.         "1": 1,
47.         "2": 1,
48.         "3": 1,
49.         "4": 2,
50.         "5": 3,
51.         "6": 4,
52.         "7": 4,
53.         "8": 1,
54.         "9": 1,
55.         "10": 2,
56.         "11": 7,
57.         "12": 8,
58.         "13": 8,
59.         "14": 1,
60.         "15": 2,
61.         "16": 3,
62.         "17": 3,
63.         "18": 4,
64.         "19": 5,
65.         "20": 5,
66.         "21": 10,
67.         "22": 2,
68.         "23": 6,
69.         "24": 1,
70.         "25": 2
71.     },
```

```

72.     "t": {
73.         "1": 0.5,
74.         "2": 1,
75.         "3": 1.5,
76.         "4": 0,
77.         "5": 0.5,
78.         "6": 0.5,
79.         "7": 0.75,
80.         "8": 2,
81.         "9": 2.5,
82.         "10": 2.5,
83.         "11": 1.5,
84.         "12": 1,
85.         "13": 1.5,
86.         "14": 4,
87.         "15": 7,
88.         "16": 3,
89.         "17": 7,
90.         "18": 3,
91.         "19": 1,
92.         "20": 6,
93.         "21": 0,
94.         "22": 6,
95.         "23": 13,
96.         "24": 2,
97.         "25": 0
98.     },
99.     "n": {
100.         "1": 0.25536998241635e1,
101.         "2": -0.44585951806696e1,
102.         "3": 0.82425886369063,
103.         "4": 0.11215007011442,
104.         "5": -0.35910933680333e-1,
105.         "6": 0.16790508518103e-1,
106.         "7": 0.32734072508724e-1,
107.         "8": 0.95571232982005,
108.         "9": -0.10003385753419e1,
109.         "10": 0.85581548803855e-1,
110.         "11": -0.25147918369616e-1,
111.         "12": -0.15202958578918e-2,
112.         "13": 0.47060682326420e-2,
113.         "14": -0.97845414174006e-1,
114.         "15": -0.48317904158760e-1,
115.         "16": 0.17841271865468,
116.         "17": 0.18173836739334e-1,
117.         "18": -0.11399068074953,

```

```

118.         "19": 0.193298966666669e-1,
119.         "20": 0.11575877401010e-2,
120.         "21": 0.15253808698116e-3,
121.         "22": -0.43688558458471e-1,
122.         "23": -0.82403190629989e-2,
123.         "24": -0.28390056949441e-1,
124.         "25": 0.14904666224681e-2
125.     },
126.     "a": {
127.         "24": 10,
128.         "25": 10
129.     },
130.     "b": {
131.         "24": 150,
132.         "25": 200
133.     },
134.     "g": {
135.         "24": 1.16,
136.         "25": 1.13
137.     },
138.     "e": {
139.         "24": 0.85,
140.         "25": 1
141.     },
142.     "n0": {
143.         "1": 12.54882924,
144.         "2": -5.46976878,
145.         "3": 3.24680487,
146.         "4": 5.54913289,
147.         "5": 11.4648996,
148.         "6": 7.59987584,
149.         "7": 9.66033239
150.     },
151.     "g0": {
152.         "4": 0.7748404445,
153.         "5": 3.3406025522,
154.         "6": 4.9705130961,
155.         "7": 9.9755537783
156.     },
157.     "last_term_ideal": 7,
158.     "last_term_residual": [
159.         7,
160.         23,
161.         25
162.     ],
163.     "phi_ideal_type": 1,

```

```

164.     "phi_residual_type": 2
165. },
166. "aux": {
167.     "reference": [
168.         "Reference Equations of State for the Thermodynamic",
169.         "Properties of Fluid Phase -Butane and Isobutane",
170.         "D. Bückner; W. Wagner"
171.     ],
172.     "delta_l_sat_approx": {
173.         "c": 1,
174.         "n": {
175.             "1": 1.97874515,
176.             "2": 0.85679951,
177.             "3": -0.341871887,
178.             "4": 0.304337558
179.         },
180.         "t": {
181.             "1": 0.345,
182.             "2": 1,
183.             "3": 1.5,
184.             "4": 3
185.         },
186.         "type": 1
187.     },
188.     "delta_v_sat_approx": {
189.         "c": 1,
190.         "n": {
191.             "1": -2.07770057,
192.             "2": -3.0836249,
193.             "3": -0.485645266,
194.             "4": -3.83167519
195.         },
196.         "t": {
197.             "1": 0.345,
198.             "2": 0.833333333,
199.             "3": 3.166666666,
200.             "4": 4.166666666
201.         },
202.         "type": 3
203.     }
204. },
205. "transport": {
206.     "thermal_conductivity": {
207.         "reference": []
208.     },
209.     "viscosity": {

```

```

210.         "reference": []
211.     },
212.     "surface_tension": {
213.         "reference": [
214.             "Mulero, A., I. Cachadina, Parra, M., Recommended Correlations for the Surface ",
215.             "    Tension of Common Fluids, J. Phys. Chem. Ref. Data 41, 043105, (2012).",
216.         ],
217.         "Tc": 425.125,
218.         "s": {
219.             "0": 0.05138
220.         },
221.         "n": {
222.             "0": 1.209
223.         },
224.         "type": 1
225.     }
226. }
227. }
228.

```

```

1. {
2.     "comp": "isobutane",
3.     "basic": {
4.         "R": 0.14305157,
5.         "MW": 58.12220,
6.         "T_star": 407.81,
7.         "rho_star": 225.5,
8.         "Tc": 407.81,
9.         "rhoc": 225.5,
10.        "Pc": 3629,
11.        "Tt": 113.73,
12.        "Pt": 0.0000219,
13.        "rhot_l": 740.339,
14.        "rhot_v": 0.00000141,
15.        "P_min": 1e-9,
16.        "P_max": 1.1e6,
17.        "rho_max": 1250.0,
18.        "T_min": 235,
19.        "T_max": 1300
20.    },
21.    "eos": {
22.        "reference": [
23.            "Reference Equations of State for the Thermodynamic",
24.            "Properties of Fluid Phase -Butane and Isobutane",

```

```
25.         "D. Bückner; W. Wagner"
26.     ],
27.     "c": {
28.         "8": 1,
29.         "9": 1,
30.         "10": 1,
31.         "11": 1,
32.         "12": 1,
33.         "13": 1,
34.         "14": 2,
35.         "15": 2,
36.         "16": 2,
37.         "17": 2,
38.         "18": 2,
39.         "19": 2,
40.         "20": 2,
41.         "21": 2,
42.         "22": 3,
43.         "23": 3
44.     },
45.     "d": {
46.         "1": 1,
47.         "2": 1,
48.         "3": 1,
49.         "4": 2,
50.         "5": 3,
51.         "6": 4,
52.         "7": 4,
53.         "8": 1,
54.         "9": 1,
55.         "10": 2,
56.         "11": 7,
57.         "12": 8,
58.         "13": 8,
59.         "14": 1,
60.         "15": 2,
61.         "16": 3,
62.         "17": 3,
63.         "18": 4,
64.         "19": 5,
65.         "20": 5,
66.         "21": 10,
67.         "22": 2,
68.         "23": 6,
69.         "24": 1,
70.         "25": 2
```

```

71.     },
72.     "t": {
73.         "1": 0.5,
74.         "2": 1,
75.         "3": 1.5,
76.         "4": 0,
77.         "5": 0.5,
78.         "6": 0.5,
79.         "7": 0.75,
80.         "8": 2,
81.         "9": 2.5,
82.         "10": 2.5,
83.         "11": 1.5,
84.         "12": 1,
85.         "13": 1.5,
86.         "14": 4,
87.         "15": 7,
88.         "16": 3,
89.         "17": 7,
90.         "18": 3,
91.         "19": 1,
92.         "20": 6,
93.         "21": 0,
94.         "22": 6,
95.         "23": 13,
96.         "24": 2,
97.         "25": 0
98.     },
99.     "n": {
100.        "1": 0.20686820727966e1,
101.        "2": -0.36400098615204e1,
102.        "3": 0.51968754427244,
103.        "4": 0.17745845870123,
104.        "5": -0.12361807851599,
105.        "6": 0.45145314010528e-1,
106.        "7": 0.30476479965980e-1,
107.        "8": 0.75508387706302,
108.        "9": -0.85885381015629,
109.        "10": 0.36324009830684e-1,
110.        "11": -0.19548799450550e-1,
111.        "12": -0.44452392904960e-2,
112.        "13": 0.46410763666460e-2,
113.        "14": -0.71444097992825e-1,
114.        "15": -0.80765060030713e-1,
115.        "16": 0.15560460945053,
116.        "17": 0.20318752160332e-2,

```

```

117.         "18": -0.10624883571689,
118.         "19": 0.39807690546305e-1,
119.         "20": 0.16371431292386e-1,
120.         "21": 0.53212200682628e-3,
121.         "22": -0.78681561156387e-2,
122.         "23": -0.30981191888963e-2,
123.         "24": -0.42276036810382e-1,
124.         "25": -0.53001044558079e-2
125.     },
126.     "a": {
127.         "24": 10,
128.         "25": 10
129.     },
130.     "b": {
131.         "24": 150,
132.         "25": 200
133.     },
134.     "g": {
135.         "24": 1.16,
136.         "25": 1.13
137.     },
138.     "e": {
139.         "24": 0.85,
140.         "25": 1
141.     },
142.     "n0": {
143.         "1": 11.60865546,
144.         "2": -5.29450411,
145.         "3": 3.05956619,
146.         "4": 4.94641014,
147.         "5": 4.09475197,
148.         "6": 15.6632824,
149.         "7": 9.73918122
150.     },
151.     "g0": {
152.         "4": 0.9512779015,
153.         "5": 2.3878958853,
154.         "6": 4.3469042691,
155.         "7": 10.3688586351
156.     },
157.     "last_term_ideal": 7,
158.     "last_term_residual": [
159.         7,
160.         23,
161.         25
162.     ],

```

```

163.     "phi_ideal_type": 1,
164.     "phi_residual_type": 2
165. },
166. "aux": {
167.     "reference": [
168.         "Reference Equations of State for the Thermodynamic",
169.         "Properties of Fluid Phase -Butane and Isobutane",
170.         "D. Bücker; W. Wagner"
171.     ],
172.     "delta_l_sat_approx": {
173.         "c": 1.0,
174.         "n": {
175.             "1": 2.04025104,
176.             "2": 0.850874089,
177.             "3": -0.479052281,
178.             "4": 0.348201252
179.         },
180.         "t": {
181.             "1": 0.355,
182.             "2": 1,
183.             "3": 1.333333333,
184.             "4": 2.3333333
185.         },
186.         "type": 1
187.     },
188.     "delta_v_sat_approx": {
189.         "c": 1,
190.         "n": {
191.             "1": -2.12933323,
192.             "2": -2.93790085,
193.             "3": -0.89441086,
194.             "4": -3.46343707
195.         },
196.         "t": {
197.             "1": 0.355,
198.             "2": 0.83333333333,
199.             "3": 3.1666666666,
200.             "4": 4.333333333
201.         },
202.         "type": 3
203.     }
204. },
205. "transport": {
206.     "thermal_conductivity": {
207.         "reference": []
208.     },

```

```
209.     "viscosity": {
210.         "reference": []
211.     },
212.     "surface_tension": {
213.         "reference": [
214.             "Mulero, A., I. Cachadina, Parra, M., Recommended Correlations for the Surface ",
215.             "    Tension of Common Fluids, J. Phys. Chem. Ref. Data 41, 043105, (2012).",
216.         ],
217.         "Tc": 407.81,
218.         "s": {
219.             "0": -0.01639,
220.             "1": 0.06121
221.         },
222.         "n": {
223.             "0": 2.102,
224.             "1": 1.304
225.         },
226.         "type": 1
227.     }
228. }
229. }
230.
```

11.2 Appendix Two: Code snippets for surrogate property package modelling

```
1. import numpy as np
2. import pandas as pd
3. from CoolProp.CoolProp import HAPropsSI, PropsSI
4. from time import perf_counter as timer
5. from multiprocessing import Pool
6.
7. # --- Constants ---
8. MW_A = PropsSI("M", "Air")
9. MW_W = PropsSI("M", "water")
10. Ttriple = PropsSI("Ttriple", "water")
11.
12. # --- Core Property Function ---
13. def compute_humid_air_properties(T, P, Y):
14.     """
15.     Computes humid air thermodynamic properties at given dry-bulb temperature (T),
16.     pressure (P), and water mole fraction (Y).
17.
18.     Returns a dictionary of properties if valid; otherwise, returns None.
19.     """
20.     try:
21.         T_w_sat_at_P = PropsSI("T", "P", P, "Q", 1, "water")
22.         if T >= T_w_sat_at_P:
23.             return None
24.
25.         Y_sat = HAPropsSI("Y", "T", T, "P", P, "RH", 1)
26.
27.         if Y <= Y_sat:
28.             MW_gas = Y * MW_W + (1 - Y) * MW_A
29.             RH = HAPropsSI("RH", "T", T, "P", P, "Y", Y)
30.             return {
31.                 "T_DB (K)": T,
32.                 "P (Pa)": P,
33.                 "x_w (mol/mol)": Y,
34.                 "x_w_gas (mol/mol)": Y,
35.                 "h_gas (J/mol)": HAPropsSI("Hha", "T", T, "P", P, "Y", Y) * MW_gas,
36.                 "s_gas (J/mol/K)": HAPropsSI("Sha", "T", T, "P", P, "Y", Y) * MW_gas,
37.                 "v_gas (m3/mol)": HAPropsSI("Vha", "T", T, "P", P, "Y", Y) * MW_gas,
38.                 "x_w_liq (mol/mol)": 1.0,
39.                 "h_liq (J/mol)": PropsSI("H", "T", T, "P", P, "water") * MW_W if T > Ttriple else
0,
40.                 "s_liq (J/mol/K)": PropsSI("S", "T", T, "P", P, "water") * MW_W if T > Ttriple
else 0,
41.                 "v_liq (m3/mol)": 1 / PropsSI("D", "T", T, "P", P, "water") * MW_W if T > Ttriple
else 0,
```

```

42.         "x_gas (mol/mol)": 1.0,
43.         "RH": RH,
44.         "T_WB (K)": HAPropsSI("Twb", "T", T, "P", P, "Y", Y)
45.     }
46.     else:
47.         x_w_gas = HAPropsSI("Y", "T", T, "P", P, "RH", 1)
48.         MW_ha = x_w_gas * MW_W + (1 - x_w_gas) * MW_A
49.         return {
50.             "T_DB (K)": T,
51.             "P (Pa)": P,
52.             "x_w (mol/mol)": Y,
53.             "x_w_gas (mol/mol)": x_w_gas,
54.             "h_gas (J/mol)": HAPropsSI("Hha", "T", T, "P", P, "RH", 1) * MW_ha,
55.             "s_gas (J/mol/K)": HAPropsSI("Sha", "T", T, "P", P, "RH", 1) * MW_ha,
56.             "v_gas (m3/mol)": HAPropsSI("Vha", "T", T, "P", P, "RH", 1) * MW_ha,
57.             "x_w_liq (mol/mol)": 1.0,
58.             "h_liq (J/mol)": PropsSI("H", "T", T, "P", P, "water") * MW_W if T > Ttriple else
0,
59.             "s_liq (J/mol/K)": PropsSI("S", "T", T, "P", P, "water") * MW_W if T > Ttriple
else 0,
60.             "v_liq (m3/mol)": 1 / PropsSI("D", "T", T, "P", P, "water") * MW_W if T > Ttriple
else 0,
61.             "x_gas (mol/mol)": (1 - Y) / (1 - x_w_gas),
62.             "RH": 1.0,
63.             "T_WB (K)": T
64.         }
65.     except Exception:
66.         return None
67.
68. # --- Input Generator ---
69. def generate_random_inputs(size, T_bounds, P_bounds, Y_bounds):
70.     """
71.     Generates a list of random (T, P, Y) input tuples within specified bounds.
72.
73.     Parameters:
74.         size (int): Number of samples
75.         T_bounds (tuple): Min and max temperature in Celsius
76.         P_bounds (tuple): Min and max pressure in kPa
77.         Y_bounds (tuple): Min and max mole fraction
78.
79.     Returns:
80.         List of tuples (T [K], P [Pa], Y [mol/mol])
81.     """
82.     T_min, T_max = T_bounds
83.     P_min, P_max = P_bounds
84.     Y_min, Y_max = Y_bounds

```

```

85.
86.     rand_T = np.random.uniform(T_min + 273.15, T_max + 273.15, size).round(6)
87.     rand_P = np.random.uniform(P_min * 1000, P_max * 1000, size).round(4)
88.     rand_Y = np.random.uniform(Y_min, Y_max, size).round(8)
89.
90.     return list(zip(rand_T, rand_P, rand_Y))
91.
92. # --- Internal Wrapper ---
93. def _get_data_point(args):
94.     """
95.     Internal helper that wraps compute_humid_air_properties and returns a DataFrame.
96.     Used for parallel execution.
97.
98.     Parameters:
99.         args (tuple): A tuple of (T, P, Y)
100.
101.     Returns:
102.         pd.DataFrame with one row or None if invalid
103.     """
104.     T, P, Y = args
105.     props = compute_humid_air_properties(T, P, Y)
106.     if props is None:
107.         return None
108.     return pd.DataFrame([props])
109.
110. # --- Dataset Generator ---
111. def generate_HAprops_dataset(
112.     project_name: str = 'pysmo_humid_air_props',
113.     size: int = 100,
114.     T_bounds = (-30, 300), # degC
115.     P_bounds = (50, 200), # kPa
116.     Y_bounds = (0.0, 0.2), # mol/mol
117.     save_csv: bool = True,
118.     serial: bool = False
119. ) -> pd.DataFrame:
120.     """
121.     Generates a dataset of humid air properties using randomized inputs.
122.
123.     Parameters:
124.         project_name (str): Output CSV filename prefix
125.         size (int): Number of data points to generate
126.         T_bounds (tuple): Temperature bounds in Celsius
127.         P_bounds (tuple): Pressure bounds in kPa
128.         Y_bounds (tuple): Water mole fraction bounds
129.         save_csv (bool): Whether to save the results to a CSV file
130.         serial (bool): Use serial (True) or parallel (False) processing

```

```

131.
132.     Returns:
133.         pd.DataFrame containing valid humid air property data
134.     """
135.     inputs = generate_random_inputs(size, T_bounds, P_bounds, Y_bounds)
136.
137.     startTime = timer()
138.
139.     if serial:
140.         results = [_get_data_point(args) for args in inputs]
141.     else:
142.         with Pool() as p:
143.             results = p.map(_get_data_point, inputs)
144.
145.     endTime = timer()
146.
147.     df_list = [r for r in results if r is not None]
148.     df_final = pd.concat(df_list, ignore_index=True)
149.
150.     print("Time taken: {:.2f}s".format(endTime - startTime))
151.     print("Valid points:", df_final.shape[0])
152.
153.     if save_csv:
154.         df_final.to_csv(f"results/{project_name}_data.csv", index=False)
155.
156.     return df_final
157.

```

11.3 Appendix Three: Refrigerant Efficiency Studies

11.3.1 The model assumptions

- Cycle analysis was conducted using a fixed evaporator temperature of 0°C.
- Minimum temperature difference of 5°C
- Set to maximise subcooling through the internal exchanger
- Modelled using isentropic, volumetric and mechanical efficiencies: 0.7, 0.8, 0.97 respectively
- Sets nominal cycle duty to 1000 kW

11.3.2 Case one: Efficiency analysis for varying sink outlet temperatures with a fixed sink inlet temperature of 15°C

The first case study aimed to investigate how different refrigerants handled an increasing sink outlet temperature from a cool inlet temperature (15°C) temperature. This case is one of the common processes the industry uses to generate hot water from cold make-up water to use on-site.




Figure 81 COP at varying outlet sink temperatures for a 15°C water sink inlet.

Figure 81 shows how COP trends against a rising water sink outlet temperature for a case where 100% of available subcooling is captured (within the other specified constraints). The results from this first study highlight how most refrigerants operate at a similar efficiency with CO₂ only performing approximately 0.1 of a COP greater than the others. The key outlier was ammonia with an operating efficiency below the other refrigerants (but still performing relatively well). The lesser efficiency is explained by the limited subcooling opportunities ammonia has, compared to the other refrigerants due to its large phase envelope.




Figure 82 T-H Curve for R32 (left) and ammonia (right) showing the limited subcooling available for ammonia.

11.3.3 Case Two: Efficiency analysis for the case of no subcooling for varying sink outlet temperatures with a fixed sink inlet temperature of 15°C




Figure 83 COP at varying outlet sink temperatures for a 15°C water sink inlet without subcooling.

The trend however is reversed for the case where subcooling is limited/available with the large ammonia phase envelope mitigating the loss of subcooling. The skinner R1234ze phase envelope causes the efficiency to decrease, as it relies on subcooling for a greater majority of its heat transfer.

11.3.4 Case Three: Pressure considerations for the different refrigerants




Figure 84 Maximum Cycle pressure

Figure 84 was produced from the same boundary conditions (15°C water inlet) and assumptions as the previous case. The figure shows most refrigerants operate at a modest pressure, between one and four bar, however, CO₂ operates at a much larger pressure (approximately 11 bar at 85°C).

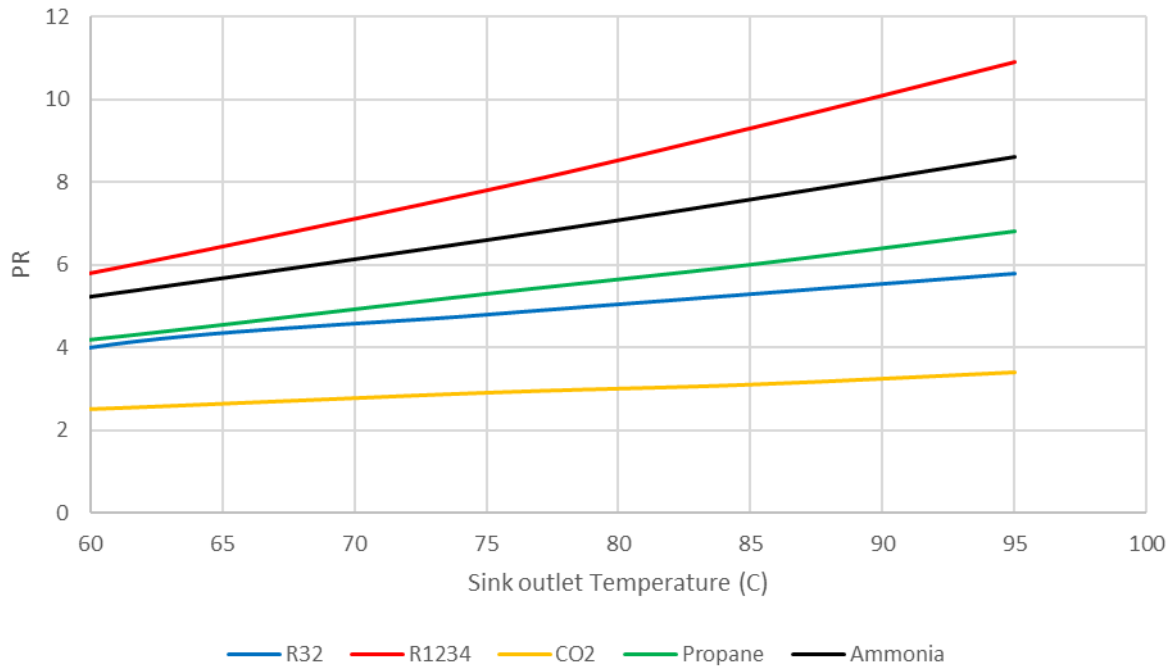


Figure 85 Pressure ratios for refrigerants at varying water outlet temperatures

During compression it is generally advantageous to keep the pressure ratio within the operating range of the manufactured device, otherwise, a series of compressors may be required. Figure 85 shows that R1234ze may require a series of compressors to achieve the target outlet temperature.

11.3.5 Case Four: Efficiency analysis for varying inlet water temperatures with a fixed water outlet condition of 85°C

It is common in industry to maintain hot water loops that have varying return temperatures; therefore, an industrial heat pump needs to perform well across a range of sink-inlet temperatures. The performance of propane, R1234 and ammonia are relatively unaffected when the sink inlet temperature increases, whereas the decrease in performance for R32 and CO₂ is more significant.

The reason for R32's performance drop and maxing out at 65°C is its lower critical point for the refrigerant limiting the efficient operable range. For CO₂, exceeding a sink temperature of 45 °C will result in a temperature cross in the gas cooler (condenser) for the given heat pump modelling constraints. The increasing inlet temperature limits the amount of heat per kg of CO₂ refrigerant that can be delivered to the water (seen in Figure 87) therefore causing more refrigerant to be pumped to maintain the same duty, which decreases the overall efficiency.




Figure 86 Efficiency for different refrigerants with varying water inlet temperatures for a fixed water outlet temperature of 85°C




Figure 87 Phase envelopes for R32 at an inlet temperature of 55°C (left), and CO2 at a water inlet temperature of 45°C (right)

11.3.6 Case Five: Simple life cycle analysis (LCA) of direct and indirect emissions for R32.

The simple LCA has the following assumptions:

- GWP AR5 for R32 is 677
- Refrigerant charge of 6kg
- Refrigerant duty of 100 kW
- COP calculated from the heat pump model
- Grid emissions factor for New Zealand was 0.12
(<https://environment.govt.nz/assets/publications/Measuring-emissions-guidance-August-2022/Summary-PDF-Measuring-emissions-guidance-August-2022.pdf>)
- Run times studied: 6000h, 3000h, 1000h
- Assumed 100% of the refrigerant charge is released after full product use
- Scope 1 emissions include refrigerant-released emissions
- Scope 2 emissions included the emissions associated with the grid

Table 27 Results for simple LCA of R32

Operational hours	Scope 1	Scope 2	Percent of the year run time
6000	1.38%	98.62%	68.49%
3000	2.72%	97.28%	34.25%
1000	7.74%	92.26%	11.42%

Table 27 shows that most of the emissions attributed to the unit come from electricity coming from the grid. This certainly raises the question of the prioritisation of refrigerant switching compared to increased cycle efficiency. Would switching to a low GWP refrigerant that operated at a worse efficiency be worse for the environment overall? Can refrigerants not be emitted into the atmosphere by being reused/disposed of properly?

11.3.7 Case Six: Summary of relevant safety standards

Most of the information comes from Proklima guidelines for the safe use of hydrocarbon refrigerants:
<https://www.irhace.org.nz/uploads/aProklima17.pdf>

Classification:

Occupancy	Description	Examples
Category A	Rooms, parts of buildings, or another location where people may sleep, where people are restricted in their movement or where the number of people present is not controlled or to which any person has access without being personally acquainted with the personal safety precautions	Hospitals, prisons, nursing homes, theatres, supermarkets, transport termini, hotels, lecture halls, dwellings, restaurants, ice rinks, passenger vehicles, etc
Category B	Rooms, parts of buildings or buildings, where only a limited number of people may be assembled, some of them being necessarily acquainted with the general safety precautions	Office buildings, laboratories, places of work, places for general manufacturing, etc
Category C	Rooms, parts of buildings, buildings where only authorised persons are granted access and which is not open to the public and; authorised persons shall be acquainted with general safety precautions of the establishment	Non public areas in supermarkets, cold stores, manufacturing facilities, refineries, manufacturing facilities e.g. for chemicals, food, ice and ice cream

The safety standard breaks down HC usage into three categories. Category A cases have small to no personal restrictions; category B limits people exposed to the HP and requires people to generally have some level of safety training. The final category C limits access to only authorised persons.

The categories are then further sectioned into how the HC HP is going to be operated. For example, in a category A case that has all or part of the system located in a human-occupied space the maximum allowable charge must not exceed 1.5kg. However, with appropriate enclosed ventilation the maximum system limit becomes $130 \times \text{LFL}(\text{kg})$ for the refrigerant (in the case of propane a max limit of 4.9kg would be specified).

For industrial hot water production most if not all hot water heat pumps would be designated category C. Sites would most likely house the compressor and receiver in a machine room or have the entire system ventilated or open to air. Therefore, the site would be required to have either a maximum charge of 25kg for the separately housed compressor or an unlimited charge for the fully ventilated/open-to-air system.

11.4 Appendix Four: All Heat Surplus-Deficit Tables for electrification retrofit

Type	Heat Exchanger Network						Process		
T _i	COND Lp Vap	Vap Bleed to Milk Preheat	COW to Milk Preheat	Flash Steam to Milk Preheat	Conc Heater	DSI to Milk Preheat	TVR E2	Evap2	Evap1
°C	ΔH _{net} kW	ΔH _{net} kW	ΔH _{net} kW	ΔH _{net} kW	ΔH _{net} kW	ΔH _{net} kW	ΔH _{net} kW	ΔH _{net} kW	ΔH _{net} kW
Upgraded Vap									
Electricity									
MP Steam					1,204	4,467	4,205		1,856
95						-3,600			614
81									
81				3,152		-26			26
78						-841			840
73				-1,189					1,188
72			329	-387					387
72		5,526	22	-26					-5,519
72			720	-846					207
69			599	-704					-7
66									
62		-1,079	916						-10
62			-26	22			-6,311	6,311	-0
57		-1,263	1,073				-52	-0	-12
57	4,118		-26	22		-1,204	2,494	-7,122	515
57									
51		-1,553	1,319				-38	92	-10
45		-1,579	1,341				-39	94	-10
45		Upto 54°C							
30				-5,70			-97	233	-24
8				-5,795			-144	346	-36
5							-20	47	-5
Cooling water	-4,118								
Bridge Sum									
Surplus	4118 kW	5526 kW				4467 kW			
Deficit		-2365 kW				-4467 kW			
Upgrade HP									
Vap Upgraded	2365 kW								
Work Required	262 kW		COP		10				
MPs Replaced	2617 kW								

Figure 88 DSI steam replacement retrofit for two-effect MVR- TVR design

Type	Heat Exchanger Network						Process		
T _i	COND Lp Vap	Vap Bleed to Milk Preheat	COW to Milk Preheat	Flash Steam to Milk Preheat	Conc Heater	DSI to Milk Preheat	TVR E2	Evap2	Evap1
°C	ΔH _{net} kW	ΔH _{net} kW	ΔH _{net} kW	ΔH _{net} kW	ΔH _{net} kW	ΔH _{net} kW	ΔH _{net} kW	ΔH _{net} kW	ΔH _{net} kW
Upgraded Vap						2,355			
Electricity						262			1,856
MV Steam					1,204	1,850	4,205		614
95									
81						-3,600			
81				3,152		-26			26
81						-841			840
78									
73				-1,189					1,188
72			329	-387					387
72		3,171	22	-26					-5,519
72			720	-846					207
69			599	-704					-7
66									
62		-1,079	916						-10
62			22				-6,311	6,311	-0
62									
57		-1,26	1,073				-52	-0	-12
57	1,763		22		-1,204		2,494	-7,122	515
57									
51			1,319				-38	92	-10
45			1,341				-39	94	-10
30				-570			-97	233	-24
30				-5,795			-144	346	-36
8									
5							-20	47	-5
Cooling water	-1,763								
Bridge Sum									
Surplus	1763 kW	3171 kW	2,609			1850 kW			
Deficit		-1,289 kW	-6,364			-4467 kW			
Upgrade HP									
Vap Upgraded	1763 kW	1665 kW							
Work Required	196 kW	185 kW		COP		10			
MPS Replaced	1959 kW	1850 kW							
	Too High	Correct							

Figure 89 DSI steam replacement retrofit for two-effect MVR-TVV design

Type	Heat Exchanger Network						Process		
T _i	COND Lp Vap	Vap Bleed to Milk Preheat	COW to Milk Preheat	Flash Steam to Milk Preheat	Conc Heater	DSI to Milk Preheat	TVR E2	Evap2	Evap1
°C	ΔH_{net} kW	ΔH_{net} kW	ΔH_{net} kW	ΔH_{net} kW	ΔH_{net} kW	ΔH_{net} kW	ΔH_{net} kW	ΔH_{net} kW	ΔH_{net} kW
Upgraded Vap						4,020			
Electricity						447			1,856
MP Steam					1,204		4,205		614
95						-3,600			
81									
81				3,152		-26			26
78						-841			840
73				-1,189					1,188
72			329	-387					387
72		1,408	22	-26					-5,519
69			720	-846					207
66			599	-704					-7
62		-1,079	916						-10
62		-26	22				-6,311	6,311	-0
62		-1,263	1,073				-52	-0	-12
57									
57	98	-26	22		-1,204		2,494	-7,122	515
51			1,319				-38	92	-10
45			1,341				-39	94	-10
30			-570				-97	233	-24
8			-4,032				-144	346	-36
5							-20	47	-5
Cooling water	-98								
Total Steam Savings: 4,467 kW									
Work: 447 kW									

Figure 90 DSI steam replacement retrofit for two-effect MVR-TV R design

Type	Heat Exchanger Network						Process		
T _i	COND Lp Vap	Vap Bleed to Milk Preheat	COW to Milk Preheat	Flash Steam to Milk Preheat	Conc Heater	DSI to Milk Preheat	TVR E2	Evap2	Evap1
°C	ΔH_{net} kW	ΔH_{net} kW	ΔH_{net} kW	ΔH_{net} kW	ΔH_{net} kW	ΔH_{net} kW	ΔH_{net} kW	ΔH_{net} kW	ΔH_{net} kW
Upgraded Vap								2,494	
Electricity					1,204	4,467	4,118	51	1,856
MP Steam						-3,600			614
				3,152		-26			26
						-841			840
				-1,189					1,188
			329	-387					387
		5,526	22	-26					-5,519
			720	-846					207
			599	-704					-7
		-1,079	916						-10
		-26	22				-3,766	3,766	-0
		-1,263	1,073				-52	-0	-12
	4,118	-26	22		-1,204			-7,122	515
57		-1,553	1,319				-38	92	-10
51		-1,579	1,341				-39	94	-10
45			-570				-97	233	-24
30			-5,795				-144	346	-36
8							-20	47	-5
5									
Cooling water	-4,118								
Bridge Sum									
Surplus	4118 kW							3766 kW	
Deficit							-3766 kW	-7122 kW	
Upgrade HP									
Vap Upgraded	3766 kW	3691 kW							
Work Required	77 kW	75 kW		COP	50				
MPS Replaced	3842 kW	3766 kW							
	Too High	Correct							

Figure 92 TVR replacement retrofit for two-effect MVR-TVV design

Type	Heat Exchanger Network						Process		
T _i	COND Lp Vap	Vap Bleed to Milk Preheat	COW to Milk Preheat	Flash Steam to Milk Preheat	Conc Heater	DSI to Milk Preheat	TVR E2	Evap2	Evap1
°C	ΔH_{net} kW	ΔH_{net} kW	ΔH_{net} kW	ΔH_{net} kW	ΔH_{net} kW	ΔH_{net} kW	ΔH_{net} kW	ΔH_{net} kW	ΔH_{net} kW
Upgraded Vap								6,185	
Electricity								128	1,856
MP Steam					1,204	4,467			614
95									
81						-3,600			
81				3,152		-26			26
78						-841			840
73				-1,189					1,188
72			329	-387					387
72		5,526	22	-26					-5,519
69			720	-846					207
66			599	-704					-7
62		-1,079	916						-10
62		-26	22						-0
57		-1,263	1,073					-0	-12
57	428	-26	22		-1,204			-7,122	515
51		-1,553	1,319					92	-10
45		-1,579	1,341					94	-10
30			-570					233	-24
8			-5,795					346	-36
5								47	-5
Cooling water	-428								

Total Steam Savings: 4,467 kW
Work: 126 kW

Figure 93 TVR replacement retrofit for two-effect MVR-TVV design

Type	Heat Exchanger Network						Process		
T _i	Cond Lp Vap	COW to Milk Preheat	Vap Bleed to Milk Preheat	Flash Vap to Milk Preheat	Conc Heat	DSI to Milk Preheat	TVR E1-4	Evap 2-3	Evap 1
°C	ΔH _{net} kW	ΔH _{net} kW	ΔH _{net} kW	ΔH _{net} kW	ΔH _{net} kW	ΔH _{net} kW	ΔH _{net} kW	ΔH _{net} kW	ΔH _{net} kW
Upgraded Vap									
Electricity									1,749
MPS					1,177	4,945	3,453		
95						-3,551			
82					3,658	-26			26
81						-631			633
79						-158			739
76						-840			844
73						-262			266
72									
22			4,542	-26	-6				-4,534
72		613		-733	-158				220
69		133		-159	-34			-6	48
69		745		-892	-192			-35	
65		22		-26	-6			-2	
65		601		-719	-155			167	
62		166	-200		-43			46	
62		22	-26		-6		-5,730	5,736	
61		701	-842		-181		-38	251	
58	3,758	22	-26				2,836	-6,596	
58		1,556	-1,869				-70	59	
51		1,315	-1,579				-59	50	
45			-732				-169	142	
28			-5,186				-195	164	
8							-29	25	
5									
Cool Water	-3,758								
Bridge Sum									
Surplus	3758 kW		4,542			4945 kW			
Deficit			-2,647						
Upgrade HP									
Vap Upgraded	2647 kW								
Work Required	294 kW		COP		10				
MPS Replaced	2941 kW								

Figure 94 DSI steam replacement retrofit for three-effect MVR-TV R design

Type	Heat Exchanger Network						Process		
T _i	Cond Lp Vap	COW to Milk Preheat	Vap Bleed to Milk Preheat	Flash Vap to Milk Preheat	Conc Heat	DSI to Milk Preheat	TVR E1-4	Evap 2-3	Evap 1
°C	ΔH _{net} kW	ΔH _{net} kW	ΔH _{net} kW	ΔH _{net} kW	ΔH _{net} kW	ΔH _{net} kW	ΔH _{net} kW	ΔH _{net} kW	ΔH _{net} kW
Upgraded Vap						2,647			
Electricity						294			1,749
MPS					1,177	2,004	3,453		
95									
82						-3,551			
81					3,658	-26			26
79						-631			633
76						-158			739
73					-840	-181			844
72					-262	-57			266
72		22	2,647			-26			-4,534
69		613				-733			220
69		133				-159		-6	48
65		745				-892		-35	
65		22				-26		-6	-2
62		601				-719		-155	167
62		16				-200		-43	46
61		22				-26		-6	
61		701				-842		-5,730	5,736
58		22				-26		-38	251
58	1,111						2,836	-6,596	
51		1,556				-801		-70	59
45		1,315						-59	50
28						-732		-169	142
8						-5,186		-195	164
5								-29	25
Cool Water	-1,111								
Bridge Sum									
Surplus	1111 kW	3025 kW	2,647			2004 kW			
Deficit		-5919 kW	-1,669						
Upgrade HP									
Vap Upgraded	1111 kW								
Work Required	123 kW	COP	10						
MPS Replaced	1234 kW								

Figure 95 DSI steam replacement retrofit for three-effect MVR-TV R design

Type	Heat Exchanger Network						Process		
T _i	Cond Lp Vap	COW to Milk Preheat	Vap Bleed to Milk Preheat	Flash Vap to Milk Preheat	Conc Heat	DSI to Milk Preheat	TVR E1-4	Evap 2-3	Evap 1
°C	ΔH _{net} kW	ΔH _{net} kW	ΔH _{net} kW	ΔH _{net} kW	ΔH _{net} kW	ΔH _{net} kW	ΔH _{net} kW	ΔH _{net} kW	ΔH _{net} kW
Upgraded Vap						3,758			
Electricity						418			1,749
MPS					1,177	770	3,453		
95						-3,551			
82				3,658		-26			26
81						-631			633
79						-158			739
76						-737			844
73				-840	-181				844
72				-262	-57				266
72			1,111	-26	-6				-4,534
69				-733	-158				220
69				-159	-34			-6	48
65		402		-892	-192			-35	
65		22		-26	-6			-2	
65		601		-719	-155			167	
62		166	-200			-43		46	
62		22	-26			-6	-5,730	5,730	
61		701	-559			-181	-38	251	
58		22					2,836	-6,596	
58		1,556					-70	59	
51		1,315					-59	50	
45							-169	142	
28		-732					-195	164	
8		-4,076					-29	25	
5									
Cool Water									
Total Steam Savings: 4,467kW									
Work: 418W									

Figure 96 DSI steam replacement retrofit for three-effect MVR-TV R design

Type	Heat Exchanger Network						Process		
T _i	Cond Lp Vap	COW to Milk Preheat	Vap Bleed to Milk Preheat	Flash Vap to Milk Preheat	Conc Heat	DSI to Milk Preheat	TVR E1-4	Evap 2-3	Evap 1
°C	ΔH _{net} kW	ΔH _{net} kW	ΔH _{net} kW	ΔH _{net} kW	ΔH _{net} kW	ΔH _{net} kW	ΔH _{net} kW	ΔH _{net} kW	ΔH _{net} kW
Upgraded Vap Electricity									1,749
MPS					1,177	4,945	2,453		
95						-3,551			
82				3,658					26
81						-631			633
79						-158			739
76									
73				-840	-181				844
72				-262	-57				266
72		22	4,542	-26	-6				-4,534
69		613		-733	-158				220
69		133		-159	-34			-6	48
65		745		-892	-192			-35	
65		22		-26	-6			-2	
62		601		-719	-155			167	
62		166	-200		-43			46	
61		22	-26		-6		-5,736	5,736	
58		701	-842		-181		-38	251	
58	3,758	22	-26				2,836	-6,596	
58		1,556	-1,869				-70	59	
51		1,315	-1,579				-59	50	
45			-732				-169	142	
28			-5,186				-195	164	
8							-29	25	
5									
Cool Water	-3,758								
Bridge Sum									
Surplus							2,836	5,736	
Deficit								-5,736	
Upgrade HP									
Vap Upgraded	2836 kW								
Work Required	58 kW								
MPS Replaced	2894 kW								
COP	5.0								

Figure 97 TVR replacement retrofit for three-effect MVR-TV design

Type	Heat Exchanger Network						Process		
T _i	Cond Lp Vap	COW to Milk Preheat	Vap Bleed to Milk Preheat	Flash Vap to Milk Preheat	Conc Heat	DSI to Milk Preheat	TVR E1-4	Evap 2-3	Evap 1
°C	ΔH _{net} kW	ΔH _{net} kW	ΔH _{net} kW	ΔH _{net} kW	ΔH _{net} kW	ΔH _{net} kW	ΔH _{net} kW	ΔH _{net} kW	ΔH _{net} kW
Upgraded								2,836	
Electricity								58	1,749
MPS					1,177	945	3,895		
95						-551			
82									
81				3,658		-26			26
79						-631			633
76						-737			739
73					-840	-181			844
72					-262	-57			266
72			22	4,542	-26	-6			-4,534
69		613			-733	-158			220
69		133			-159	-34		-6	48
65		745			-892	-192		-35	
65		22			-26	-6		-2	
62		601			-719	-155			167
62		166	-200			-43			46
61		22	-26			-6	-2,842	2,842	
58		701	-842			-181	-38	251	
58	3,758	22	-26					-6,596	
51		1,556	-1,869				-70	59	
45		1,315	-1,579				-59	50	
28		-732					-169	142	
8		-5,186					-195	164	
5							-29	25	
Cool Water	-3,758								
Bridge Sum									
Surplus	3758 kW						2,842	2,842	
Deficit							-2,842	-2,842	
Upgrade HP									
Vap Upgraded	2842 kW	2785 kW							
Work Required	58 kW	57 kW							
MPS Replaced	2900 kW	2842 kW							
COP	Too high								
	50								

Figure 98 TVR replacement retrofit for three-effect MVR-TV design

Type	Heat Exchanger Network						Process		
T _i	Cond Lp Vap	COW to Milk Preheat	Vap Bleed to Milk Preheat	Flash Vap to Milk Preheat	Conc Heat	DSI to Milk Preheat	TVR E1-4	Evap 2-3	Evap 1
°C	ΔH _{net} kW	ΔH _{net} kW	ΔH _{net} kW	ΔH _{net} kW	ΔH _{net} kW	ΔH _{net} kW	ΔH _{net} kW	ΔH _{net} kW	ΔH _{net} kW
Upgraded Vap								5,678	
Electricity								115	1,749
MPS					1,177	4,945			
95						-3,551			
82				3,658		-26			26
81						-631			633
79						-737			739
76					-158				
73				-840	-181				844
72				-262	-57				266
72		22	4,542	-26	-6				-4,534
69		613		-733	-158				220
69		133		-159	-34			-6	48
69		745		-892	-192			-36	
65		22		-26	-6			-2	
65		601		-719	-155			167	
62		166	-200		-43			46	
62		22	-26		-6				
61		701	-842		-181			251	
58	973	22	-26					-6,596	
58		1,556	-1,869					59	
51		1,315	-1,579					50	
45		-732						142	
28		-5,186						164	
8								25	
5									
Cool Water	-973								
Total Steam Savings: 3453kW									
Work: 115W									

Figure 99 TVR replacement retrofit for three-effect MVR-TV design

		Heat Exchanger Network				Process	
T_i	Cond Lp	COW to Milk Preheat	Vap Bleed to Milk Preheat	Conc Preheat	DSI to Milk Preheat	Evap 1-4	TVR E1-4
$^{\circ}\text{C}$	ΔH_{net} kW	ΔH_{net} kW	ΔH_{net} kW	ΔH_{net} kW	ΔH_{net} kW	ΔH_{net} kW	ΔH_{net} kW
Upgraded Vap							
Electricity							
MP Steam				1,382	6,569		13,900
95					-4,204	4,382	
79				-338	-1,577	1,643	
73			6,577	-6	-26	22,115	-28,669
73				-164	-762	794	-148
70			-517	-111		539	-100
68			-719	-154		604	-139
65		124	-137	-29		115	-27
65		865	-954	-205		605	-185
61		1,564	-1,725	-370		734	-334
55	14,725	24	-26	-6		-31,610	16,780
55							
51	835		-921	Upto 52°C		5	-76
45			-1,579			9	-131
25		-473				32	-442
8		-4,370				26	-363
5						5	-65
Cooling Water	-14,725						
Bridge Sum							
Surplus	14725 kW		6,577		6,569		
Deficit			-1,809				
Upgrade HP							
Vap Upgraded	1809 kW						
Work Required	201 kW		COP	10			
MPS Replaced	2010 kW						

Figure 100 DSI steam replacement retrofit for four-effect TVR design

T _i	Heat Exchanger Network					Process	
	Cond Lp	COW to Milk Preheat	Vap Bleed to Milk Preheat	Conc Preheat	DSI to Milk Preheat	Evap 1-4	TVR E1-4
	ΔH_{net} kW	ΔH_{net} kW	ΔH_{net} kW	ΔH_{net} kW	ΔH_{net} kW	ΔH_{net} kW	ΔH_{net} kW
°C							
Upgraded Vap					2,798		
Electricity					311		
MP Steam				1,382	3,460		13,900
95							
					-4,204	4,382	
79							
				-338	-1,577	1,643	
73							
			3,779	-6	-26	22,115	-28,669
73							
				-164	-762	794	-148
70							
			-517	-111		539	-100
68							
			-719	-154		604	-139
65							
		124	-137	-29		115	-27
65							
		865	-954	-205		605	-185
61							
		1,564	-1,452	-370		734	-334
55							
	11,927	24		-6		-31,610	16,780
55							
		835				5	-76
51							
		1,431				9	-131
45							
		-473				32	-442
25							
						26	-363
8							
						5	-65
5							
Cooling Water	-11,927						

Steam Savings: **3109KW**
Work: **311kW**

Figure 102 DSI steam replacement retrofit for four-effect TVR design

		Heat Exchanger Network					Process	
T_i	Cond Lp	COW to Milk Preheat	Vap Bleed to Milk Preheat	Conc Preheat	DSI to Milk Preheat	TVR E1-4	Evap 1-4	
$^{\circ}\text{C}$	ΔH_{net} kW	ΔH_{net} kW	ΔH_{net} kW	ΔH_{net} kW	ΔH_{net} kW	ΔH_{net} kW	ΔH_{net} kW	
Upgraded Vap								
Electricity								
MP Steam				1,382	6,569	13,900		
95					-4,204		4,382	
79							1,643	
73				-338	-1,577			
73			6,577	-6	-26	-28,669	22,115	
73								
70				-164	-762		794	
68			-517	-111			539	
65			-719	-154		139	604	
65		124	-137	-29		-27	115	
61		865	-954	-205		-185	605	
61		1,564	-1,725	-370		-334	734	
55	14,725	24	-26	-6		16,780	-31,610	
55		835	-921			-76	5	
51		1,431	-1,579			-131	9	
45		-473				-442	32	
25		-4,370				-363	26	
8						-65	5	
5								
Cooling Water	-14,725							
Bridge Sum								
Surplus						16,780 kW	22,115	
Deficit							22,115	
Upgrade HP								
Vap Upgraded	16780 kW							
Work Required	1864 kW		COP	10				
MPS Replaced	18644 kW							

Figure 103 TVR replacement retrofit for four-effect TVR design

		Heat Exchanger Network				Process	
T _i	Cond Lp	COW to Milk Preheat	Vap Bleed to Milk Preheat	Conc Preheat	DSI to Milk Preheat	TVR E1-4	Evap 1-4
°C	ΔH_{net} kW	ΔH_{net} kW	ΔH_{net} kW	ΔH_{net} kW	ΔH_{net} kW	ΔH_{net} kW	ΔH_{net} kW
Upgraded Vap							16,780
Electricity							1,864
MP Steam				1,382	6,569	12,031	
95					-4,204		4,382
79							
73				-338	-1,577		1,643
73			6,577	-6	-26	-10,025	3,471
73							
70				-164	-762	-148	794
68			-517	-111		-100	539
65			-719	-154		-139	604
65		124	-137	-29		-27	115
61		865	-954	-205		-185	605
55	14,725	1,564	-1,725	-370		-334	734
55		24	-26	-6			-31,610
55		835	-921			-76	5
51		1,431	-1,579			-131	9
45		-473				-442	32
25		-4,370				-363	26
8						-65	5
5							
Cooling Water	-14,725						
Bridge Sum							
Surplus	14725 kW						3,471
Deficit							3,471
Upgrade HP							
Vap Upgraded	3471 kW	3124 kW					
Work Required	386 kW	347 kW	COP	10			
MPS Replaced	3857 kW	3471 kW					
	Too High	Correct					

Figure 104 TVR replacement retrofit for four-effect TVR design

	Heat Exchanger Network					Process	
T _i	Cond Lp	COW to Milk Preheat	Vap Bleed to Milk Preheat	Conc Preheat	DSI to Milk Preheat	TVR E1-4	Evap 1-4
°C	ΔH _{net} kW	ΔH _{net} kW	ΔH _{net} kW	ΔH _{net} kW	ΔH _{net} kW	ΔH _{net} kW	ΔH _{net} kW
Upgraded Vap							20,251
Electricity							2,212
MP Steam				1,382	6,569	8,565	
95							
79					-4,204		4,382
73				-338	-1,577		1,643
73			6,577	-6	-26	-6,554	
73				-164	-762	-148	794
70							
68				-517	-111	-100	539
65							
65		124	-137	-29		-27	115
65		865	-954	-205		-185	605
61		1,564	-1,725	-370		-334	734
55	11,601	24	-26	-6			-31,610
55		835	-921			-76	5
51		1,431	-1,579			-131	9
45							
45			-473			-442	32
25							
25			-4,370			-363	26
8							
8						-65	5
5							
Cooling Water	-11,601						

Steam Saving: **5335kW**
Work: **2212kW**

Figure 105 TVR replacement retrofit for four-effect TVR design

11.5 Appendix Five: Evidence of contributions to published articles



Co-Authorship Form

School of Graduate Research
The University of Waikato
Private Bag 3105
Hamilton 3240, New Zealand
Phone +64 7 838 5096
Email: SGR@waikato.ac.nz
Website: <http://www.waikato.ac.nz/students/research>

This form is to accompany the submission of any PhD that contains research reported in published or unpublished co-authored work. **Please include one copy of this form for each co-authored work.** Completed forms should be included in your appendices for all the copies of your thesis submitted for examination and library deposit (including digital deposit).

Chapter 7 - Process integration and electrification for efficient milk evaporation systems DOI: <https://doi.org/10.1016/j.energy.2022.124885>

Nature of contribution by PhD candidate	Writing – original draft, Visualization, Methodology, Investigation, Formal analysis, Conceptualization.
Extent of contribution by PhD candidate (%)	85


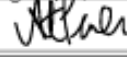
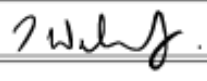
CO-AUTHORS

Name	Nature of Contribution
Lana Kong	Writing – original draft, Visualization, Formal analysis
Alyssa Pineda	Formal analysis
Timothy Walmsley	Writing – review & editing, Supervision, Software, Funding acquisition, Data curation, Conceptualization.

Certification by Co-Authors

The undersigned hereby certify that:

❖ the above statement correctly reflects the nature and extent of the PhD candidate's contribution to this work, and the nature of the contribution of each of the co-authors; and ❖ that the candidate wrote all or the majority of the text.

Name	Signature	Date
Lana Kong		25/08/2025
Alyssa Pineda		25/08/2025
Timothy Walmsley		25/08/2025

July 2015



Co-Authorship Form

School of Graduate Research
The University of Waikato
Private Bag 3105
Hamilton 3240, New Zealand
Phone +64 7 838 5096
Email: SGR@waikato.ac.nz
Website: <http://www.waikato.ac.nz/students/research>

This form is to accompany the submission of any PhD that contains research reported in published or unpublished co-authored work. **Please include one copy of this form for each co-authored work.** Completed forms should be included in your appendices for all the copies of your thesis submitted for examination and library deposit (including digital deposit).

Chapter 8 - Process Integration and Electrification for retrofit: Case studies of milk evaporator systems DOI: <https://doi.org/10.1016/j.csite.2024.104601>

Nature of contribution by PhD candidate

Writing – original draft, Visualization, Methodology, Investigation, Formal analysis, Conceptualization.

Extent of contribution by PhD candidate (%)

85%



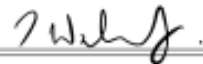
CO-AUTHORS

Name	Nature of Contribution
Lana Kong	Writing – original draft, Visualization.
Florian Schlosser	Writing – original draft, Visualization, Investigation, Formal analysis.
Timothy Walmsley	Writing – review & editing, Supervision, Software, Funding acquisition, Data curation, Conceptualization.

Certification by Co-Authors

The undersigned hereby certify that:

❖ the above statement correctly reflects the nature and extent of the PhD candidate's contribution to this work, and the nature of the contribution of each of the co-authors; and ❖ that the candidate wrote all or the majority of the text.

Name	Signature	Date
Lana Kong		25/08/2025
Florian Schlosser		25.08.2025
Timothy Walmsley		25/08/2025

July 2015