



Trinity College Dublin

Coláiste na Tríonóide, Baile Átha Cliath

The University of Dublin

Vapour-Compression Assisted Thermal Management of High-Powered Electronics

Fazeel Mohammed Naduvilakath Mohammed

Supervisor: Dr. Anthony James Robinson

A thesis submitted in partial fulfilment for the degree of

Doctor of Philosophy

in Mechanical and Manufacturing Engineering

Department of Mechanical, Manufacturing & Biomedical Engineering,

Trinity College Dublin, The University of Dublin, Ireland

2026

Declaration of Authorship

I declare that this thesis has not been submitted as an exercise for a degree at this or any other university and it is entirely my own work. I confirm that the work has been completed in full compliance with relevant university policies, including the Academic Integrity policy, the Trinity Policy on Good Research Practice, and the guidance on AI and GenAI in Teaching, Learning, Assessment and Research.

I agree to deposit this thesis in the University's open access institutional repository or allow the library to do so on my behalf, subject to Irish Copyright Legislation and Trinity College Library conditions of use and acknowledgement.

I consent to the examiner retaining a copy of the thesis beyond the examining period, should they so wish (EU GDPR May 2018).

A handwritten signature in black ink, appearing to read 'Fazeel', with a horizontal line drawn through it.

Fazeel Mohammed N. M.

Abstract

The research presented in this thesis investigates the evolving thermal-hydraulic and energy challenges associated with cooling modern high power Central Processing Units (CPUs), culminating in a unified modelling, optimisation, and system integration framework for next generation active cooling and heat recovery concepts. Through a sequence of experimental, theoretical, and mathematical studies, the work develops a systematic understanding of the interplay between heat transfer performance, fluid-mover penalties, and system level optimisation strategies for hybrid-liquid, vapour-compression refrigeration and heat pump systems.

The first stage of the research characterised the cooling performance of traditional Fan-Fin Air-Cooling (FFAC) and Hybrid Liquid-Air (HLA) systems to characterise their performance and establish a realistic baseline against which all subsequent active cooling concepts could be evaluated. This study quantified the full thermal-hydraulic envelope of both systems, demonstrating the clear limitations of FFAC and the superior performance of the HLA approach, while also showing how coordinated fluid-mover control can minimise power consumption penalties. In terms of pre-emptive engineering for the escalating heat fluxes of CPUs, a major limitation was found to be that of the ambient temperature heat sink.

Building on this baseline, the second phase investigated Vapour-Compression Refrigeration (VCR) assisted liquid-cooling, first experimentally and then through validated full system modelling. Experimental results confirmed that miniature VCR-assisted liquid-cooling can deliver sub-ambient coolant temperatures and significantly extend CPU cooling capability while maintaining stable operation across wide ranges of heat load, flow rate, and compressor speed. However, the dominant role of compressor power in governing fluid-mover energy consumption penalties emphasised the necessity of coordinated control strategies to achieve favourable system Coefficient of Performance (COP) levels. To address this, a comprehensive steady-state mathematical model of the VCR-assisted liquid-cooling system was developed, validated against experiments, and integrated with a Design of Experiments (DOE) methodology. This modelling framework enabled rigorous parameter sensitivity analysis and system-level optimisation,

demonstrating that properly selected combinations of coolant flow rate, compressor speed, and heat exchanger sizing can achieve target CPU die temperatures with mathematically minimized net power consumption, achieving up to COP = 6.5.

The final stage of the research expanded the vapour-compression modelling framework to a dual-purpose refrigeration and heat pump system capable of simultaneously maintaining suitably low liquid coolant temperatures for CPU cooling while upgrading waste heat to $\geq 80^{\circ}\text{C}$, making it suitable as a heat recovery system for Direct Air Capture (DAC) applications. A new heat recovery mathematical system model, including a high-pressure ratio compressor and a water-cooled condenser heat exchanger, was developed and embedded into a DOE-driven optimisation workflow. Case studies demonstrated that, with appropriate sizing of the evaporator and condenser heat exchangers and coordinated compressor-coolant control, the miniature system could jointly satisfy CPU thermal constraints and DAC-suitable heat recovery temperatures, achieving COPs between 3.8 and 5.9 under optimised conditions. This work represents the first demonstration of an integrated miniature vapour-compression system simultaneously optimised for CPU cooling and high-temperature heat recovery.

Overall, this research provides a cohesive and experimentally grounded framework for analysing, comparing, and optimising advanced CPU cooling architectures. By establishing baseline liquid-air performance, developing validated VCR modelling tools, and extending these to heat recovery applications, the research delivers new insight into the fundamental coupling between heat transfer, thermodynamics, and fluid-mover penalties in hybrid cooling systems. The outcomes offer a pathway toward high performance, energy efficient, and multifunctional thermal management solutions tailored to the rapidly evolving demands of modern high performance computing components and systems.

Publications and Conferences

Journal Publications

- [1] F.M. Naduvilakath-Mohammed, R. Jenkins, G. Byrne, A.J. Robinson, Closed loop liquid cooling of high-powered CPUs: A case study on cooling performance and energy optimization, *Case Studies in Thermal Engineering* 50 (2023).
- [2] F.M. Naduvilakath-Mohammed, M. Lebon, G. Byrne, A.J. Robinson, An experimental investigation of Vapour-Compression refrigeration cooling and energy performance for CPU thermal management, *Case Studies in Thermal Engineering* 54 (2024).

Conference Papers

- [1] F.M. Naduvilakath-Mohammed, M. Lebon, G. Byrne, G. Murray, S. Battagoli, A.J. Robinson, Thermal Performance of Miniature Vapour-Compression Refrigeration System for CPU Thermal Management, in: *2022 28th International Workshop on Thermal Investigations of ICs and Systems (THERMINIC)*, 2022: pp. 1–4.
- [2] F.M. Naduvilakath-Mohammed, M. Lebon, A.J. Robinson, Numerical modelling of a hybrid Vapour-Compression refrigeration assisted closed loop liquid cooling system for high-performance computing systems, *J. Phys.: Conf. Ser.* 2766 (2024).

Acknowledgements

I would like to express my deepest gratitude to my supervisor, Dr. Anthony James Robinson (Tony), Associate Professor in the Department of Mechanical, Manufacturing and Biomedical Engineering at Trinity College Dublin, for his unwavering guidance, encouragement, and mentorship throughout this research. His expertise and support have been instrumental in my development as a researcher, and it has been a privilege to work under his supervision.

I also sincerely thank Mr. Gerard Byrne (Gerry), Senior Experimental Officer, for his invaluable technical support throughout the experimental work, as well as the workshop staff for their professionalism. I am equally grateful to Trinity College Dublin for providing an exceptional research environment and to Nexalus and its team for their research funding, technical support, and the invaluable opportunity to undertake a research internship, which provided valuable industry experience and exposure to real-world engineering challenges..

I wish to express my deepest and most heartfelt gratitude to my parents, Uppa and Umma, for their unconditional love, endless sacrifices, unwavering belief in me, and constant prayers. Leaving home to pursue my master's degree in Ireland was the first time I had ever been away from my family after 26 years of my life, and the early days were incredibly difficult. The homesickness was overwhelming, but no matter the time of day or night, they were always just a phone call away-listening to my worries, celebrating my smallest achievements, and reminding me that I was never alone. Their love, strength, and faith carried me through every challenge, including the uncertainty of the COVID-19 pandemic. Everything I have achieved today is because they first believed in me. This achievement is as much theirs as it is mine.

To my little sister, Pathu, thank you for being my constant source of strength. You stood by me like a shadow, especially through some of the most difficult personal moments of my life. Even from miles away, you always found ways to make my life easier, safer, and brighter. Your selflessness, love, and unwavering support gave me the courage to keep moving forward when I needed it most. You

are truly the best sibling anyone could ever ask for, and I will always be grateful to have you in my life.

I would also like to acknowledge my young nephew, Rayyan, whose innocent laughter, cheerful giggles, and heartwarming smiles during our video calls brought immense joy and comfort during some of the most stressful moments of my PhD thesis writing. Seeing his excitement whenever he saw his uncle on the screen provided a welcome escape from the pressures of research, helping me find calm, happiness, and perspective. Though he may be too young to understand it now, his presence became a meaningful source of strength and positivity throughout this journey.

I am also deeply grateful to my wife, Sumi, who joined my life in the latter half of my PhD journey and became a constant source of comfort and support. Her patience, understanding, and unwavering belief in me - especially during the final stages and submission of this thesis - gave me the strength to persevere. This work would not have been possible without her love and calm strength.

I am equally thankful to my friends, research colleagues, and well-wishers for their encouragement and companionship throughout this journey. I gratefully acknowledge the financial support provided by the Irish Research Council (IRC) through the IRC-EPS Award EPSPG/2022/347, and by the CONNECT Research Centre, supported by Science Foundation Ireland (SFI) under Grant Number 13/RC/2077_P2.

Above all, I acknowledge the strength and guidance beyond myself that sustained me throughout this journey, offering resilience during moments of uncertainty and the calm resolve to persevere through challenges.

The past few years have been a period of many firsts, marked by silent struggles, personal growth, and resilience. This journey has been both challenging and deeply rewarding, and looking back, I am grateful for every experience that shaped me into a stronger and more confident individual. This thesis stands as a reflection of the journey undertaken, the lessons learned, and the perseverance developed along the way.

Contents

Declaration of Authorship	iii
Abstract.....	iv
Publications and Conferences.....	vi
Acknowledgements	vii
Contents.....	ix
List of Figures.....	xiv
List of Tables.....	xxi
Nomenclature.....	xxiii
1. Introduction	1
1.1. Background and Motivation	1
1.1.1. Data centre Energy Trends and Thermal Challenges	1
1.1.2. Air-Cooling: Capabilities and Constraints	5
1.1.3. Liquid-Cooling: Advantages and Emerging Limitations	7
1.1.4. Coolant Temperature as a Limiting Factor	9
1.1.5. The Role of Active Cooling in Future Data Centres.....	10
1.1.6. Waste-Heat Recovery as an Emerging Driver for Advanced Cooling Architectures.....	12
1.1.7. Towards Integrated Cooling and Heat-Recovery Systems.....	14
1.2. Research Objectives.....	15
1.3. Potential Research Impact.....	17
1.4. Thesis Structure.....	17
1.5. References	20
2. Closed Loop Liquid-Cooling of High-Powered CPUs: A Case Study on Cooling Performance and Energy Optimization	27
2.1. Introduction	28

2.2. Experimental Apparatus and Data Reduction	33
2.2.1. Thermal Test Vehicle with Liquid-Cooled Cold Plate and Remote Air-Cooled Heat Exchanger.....	33
2.2.2. Fan-Fin CPU Cooler	39
2.2.3. Uncertainty Analysis	41
2.3. Results and Discussion.....	42
2.3.1. Thermal Characterisation.....	42
2.3.2. Liquid Flow Characterisation - HLA System.....	49
2.3.3. Fluid Characterisation - Air Side	53
2.3.4. Cooling Energy Optimization	56
2.4. Conclusions	60
2.5. References.....	61
3. An Experimental Investigation of Vapour-Compression Refrigeration Cooling and Energy Performance for CPU Thermal Management	66
3.1. Introduction	67
3.2. Experimental Apparatus and Data Reduction	70
3.2.1. Thermal Test Vehicle with Liquid-Cooled Cold plate and Vapour-Compression Refrigeration	70
3.2.2. Uncertainty Analysis	78
3.3. Results and Discussion.....	79
3.3.1. Effect of Heat Load and Compressor Speed	79
3.3.2. Effect of Primary (Liquid) and Tertiary (Air) Flow Rates.....	85
3.3.3. Power optimization of VCR system.....	89
3.3.4. Comparison of VCR and HLA Cooling Systems	92
3.4. Conclusions	93
3.5. References.....	95
4. Modelling and Optimization of Vapour-Compression Refrigerated Liquid-Cooling of CPUs	99

4.1. Introduction	100
4.2. Methodology	104
4.2.1. Model Overview	104
4.2.2. Refrigerant Cooling Loop	106
4.2.3. Liquid Cooling Loop	114
4.2.4. System Energetics	114
4.3. Calculation Algorithm	115
4.4. Design of Experiments	117
4.5. Results and Discussions	118
4.5.1. Model Validation	118
4.5.2. Validation of the DOE-Based Optimization Framework	120
4.5.3. Parameter Sensitivity Analysis	122
4.5.4. Design Optimization Case Study	125
4.6. Conclusions	134
4.7. References	136
5. Modelling and Optimization of a Dual-Purpose Vapour-Compression System for CPU Liquid-Cooling and High Temperature Heat Recovery	139
5.1. Introduction	140
5.2. Experimental Apparatus	143
5.3. Mathematical Modelling	147
5.3.1. Model Overview	147
5.3.2. Refrigerant Cooling Loop	149
5.3.3 Liquid-Cooling Loop	154
5.3.4. System Energetics: Heat Recovery Model	154
5.3.5. Model Validation	155
5.4. Design of Experiments	157
5.5. Results and Discussions	158

5.5.1. Parameter Sensitivity Analysis.....	158
5.5.2. Design Optimization Case Study	162
5.6. Conclusions	173
5.7. References.....	175
6. Conclusions and Future Work	181
6.1. Conclusions	181
6.2. Future Work	183
Appendix A: Empirical Correlations and Heat Transfer Modelling.....	186
A.1. Air-Side Heat Transfer Correlation	186
A.2. Single-Phase Refrigerant Flow	186
A.3. Condensation Region	187
A.4. Evaporation Region	188
A.5. Liquid-Side Heat Transfer in the Primary Loop.....	189
A.6. Determination of Log-Mean Temperature Differences (LMTDs).....	189
A.6.1. Condenser.....	190
A.6.2. Evaporator.....	190
A.7. Effectiveness-NTU Relations.....	191
A.8. References	192
Appendix B: Calibrations	194
B.1. Thermocouple Calibration	194
B.2. Differential Pressure Sensor Calibration.....	197
B.3. Liquid Flow Meter Calibration	198
B.4. Calibration Certificates	199
Appendix C: Properties of the Freecor - Water Coolant Mixture	202
C.1. Boiling Point (°C).....	202
C.2. Freezing Point (°C).....	202
C.3. Density (kg/m ³).....	202

C.4. Thermal conductivity (W/mK)	203
C.5. Kinematic viscosity (cSt = mm ² /s).....	204
C.6. Specific heat (kJ/kgK).....	204
Appendix D: Experimental Apparatus.....	206
D.1. Thermal Test Vehicle.....	206
D.2. Cold Plate	207
D.3. Air Cooled Heat Exchanger (ACHE).....	207
D.4. Fan-Fin Air-Cooler (FFAC)	208
D.5. Aspen VCR System.....	208
D.6. RIGID VCR System	209

List of Figures

Figure 1.1. Forecasted increase in European data centre electricity consumption and its share of total power demand between 2023 and 2030	2
Figure 1.2. Historical trends in microprocessor technology	3
Figure 1.3. Typical server node illustrating the principal heat-generating components within modern computing hardware, including the CPU, DIMMs, and Voltage Regulator Modules (VRMs)	4
Figure 1.4. Simplified Direct Air Capture (DAC) process illustrating atmospheric CO ₂ capture by a sorbent and subsequent thermally driven regeneration to release concentrated CO ₂ for storage or utilisation.	13
Figure 2.1. Thermal resistance network of a CPU package.	28
Figure 2.2. Wireframe diagram of the liquid-cooling system using remote air-cooled heat exchanger.	33
Figure 2.3. (a) Schematic of the Thermal Test Vehicle (TTV) with mounted liquid-cooled cold plate, and (b) illustration of heated meter bar and the resulting linear Temperature vs. Distance plot.	36
Figure 2.4. Wireframe diagram of the impingement flow Fan-Fin Air-Cooler (FFAC) system.	39
Figure 2.5. Thermal resistance of FFAC system with varied heat loads; T _{ambient} = 35°C.	43
Figure 2.6. Expected cooling performance of a FFAC at maximum fan Speed with varying heat loads: T _{ambient} = 35°C and R _{TIM1} = 0.1 K-cm ² /W.	44
Figure 2.7. Power Consumption (PC) & Coefficient of Performance (COP) versus heat load for FFAC system.	45

Figure 2.8. Thermal resistance of HLA cooling system with varied heat loads; $T_{\text{ambient}} = 35^{\circ}\text{C}$.	46
Figure 2.9. Expected cooling performance of HLA cooling system with varying heat loads for a nominal flow rate of $Q_{\text{liq}} = 4 \text{ L/min}$ and fan speed of $v = 3 \text{ m/s}$; $T_{\text{ambient}} = 35^{\circ}\text{C}$ and $R_{\text{TIM1}} = 0.1 \text{ K-Cm}^2/\text{W}$.	47
Figure 2.10. Pressure drop curves of the HLA system for varying heat loads.	48
Figure 2.11. Power Consumption & Coefficient of Performance profile of the HLA system.	49
Figure 2.12. Thermal resistance of HLA system for varying liquid flow rates measured at constant heat flux and air velocity; $T_{\text{ambient}} = 35^{\circ}\text{C}$.	50
Figure 2.13. Expected cooling performance curves of HLA system for varying liquid flow rates at constant heat flux and air velocity; $T_{\text{ambient}} = 35^{\circ}\text{C}$ and $R_{\text{TIM1}} = 0.1 \text{ K-cm}^2/\text{W}$.	51
Figure 2.14. Pressure drop of HLA system for varying liquid flow rates at constant heat load.	52
Figure 2.15. Coefficient of Performance of HLA system for three different heat fluxes with varying liquid flow rates and constant air velocity.	53
Figure 2.16. Thermal resistance of HLA system for varying air velocities at constant heat flux and liquid flow rate; $T_{\text{ambient}} = 35^{\circ}\text{C}$.	54
Figure 2.17. Temperature profile of the HLA system for varying air velocities measured at constant heat flux and liquid flow rate; $T_{\text{ambient}} = 35^{\circ}\text{C}$ and $R_{\text{TIM1}} = 0.1 \text{ K-cm}^2/\text{W}$.	55
Figure 2.18. Power Consumption and Coefficient of Performance of the HLA System for varying air velocities at constant heat flux and liquid flow rate.	56

Figure 2.19. HLA Power Consumption (PC) surface plots for maintaining the die temperature at (a) $T_{\text{die}} = 75^{\circ}\text{C}$ and (b) $T_{\text{die}} = 55^{\circ}\text{C}$, by varying liquid flow rate and air velocity.	57
Figure 2.20. HLA Coefficient of Performance (COP) surface plots for maintaining the die temperature at (a) $T_{\text{die}} = 75^{\circ}\text{C}$ and (b) $T_{\text{die}} = 55^{\circ}\text{C}$, by varying liquid flow rate and air velocity.	59
Figure 3.1. Schematic of the Thermal Test Vehicle (TTV) with mounted liquid-cooled cold plate.	70
Figure 3.2. Illustration of heated meter bar and the resulting linear Temperature vs. Distance plot.	71
Figure 3.3. Wireframe diagram of the miniature VCR-assisted liquid-cooling system.	73
Figure 3.4. Wireframe diagram of the instrumented HLA cooling system.	77
Figure 3.5. Thermal indicator for increasing CPU heat loads at $Q_{\text{liq}} = 4 \text{ L/min}$ and $v = 3.7 \text{ m/s}$ for compressor speeds of (a) $S = 6500 \text{ rpm}$, and (b) $S = 2540 \text{ rpm}$.	80
Figure 3.6. Cooling performance of VCR-assisted liquid-cooled system for varying heat load at $Q_{\text{liq}} = 4 \text{ L/min}$ and $v = 3.7 \text{ m/s}$ for compressor speeds of $S = 6500 \text{ rpm}$, and $S = 2540 \text{ rpm}$.	81
Figure 3.7. P-h diagram of VCR-assisted liquid-cooled system for varying heat load at $Q_{\text{liq}} = 4 \text{ L/min}$ and $v = 3.7 \text{ m/s}$ for compressor speed of (a) $S = 6500 \text{ rpm}$, and (b) $S = 2540 \text{ rpm}$.	82
Figure 3.8. Power penalty plots along with associated COP for low and high CPU power and compressor speeds; $Q_{\text{liq}} = 4 \text{ L/min}$ and $v = 3.7 \text{ m/s}$.	84

Figure 3.9. Thermal indicator for varying (a) primary (liquid) flow rate and (b) tertiary (air) speed for maximum tested heat flux of $q'' = 75 \text{ W/cm}^2$ and compressor speed $S = 6500 \text{ rpm}$.	86
Figure 3.10. Cooling performance of VCR-assisted liquid-cooled system for varying primary (liquid) flow rate, and varying tertiary (air) speed for maximum tested heat flux of $q'' = 75 \text{ W/cm}^2$ and compressor speed $S = 6500 \text{ rpm}$ and 2540 rpm .	87
Figure 3.11. Power penalty plots along with associated COP for varying primary (liquid) flow rate, and varying tertiary (air) speed for maximum tested heat load (75 w/cm^2) and compressor speed (6500 rpm).	88
Figure 3.12. P-h diagram of VCR liquid-cooled system for varying condenser air velocity with fixed compressor speed, liquid flow rate and heat load.	89
Figure 3.13. (a) VCR Power Consumption (PC), and (b) COP surface plots for maintaining a die temperature of $T_{\text{die}} = 55 \text{ }^\circ\text{C}$ for varying heat loads by varying liquid flow rate and compressor speed for $T_{\text{ambient}} = 20^\circ\text{C}$.	91
Figure 3.14. Cooling performance of FFAC, HLA and VCR cooling systems for varying heat load.	92
Figure 4.1. (a) Schematic of the modelled system, (b) P-h diagram schematic of a Vapour-Compression cycle with sub-cooling and superheating, and (c) description of the operating points.	105
Figure 4.2. Schematic representation and geometric specifications of the cross-flow air-cooled condenser.	109
Figure 4.3. Schematic of the plate heat exchanger (evaporator) and key geometric specifications.	111
Figure 4.4. Flowchart of the solution algorithm of the mathematical model.	116

Figure 4.5. Comparison of predicted and measured system temperatures over a wide range of heat source powers, compressor speeds, air speeds, and liquid coolant flow rates.	119
Figure 4.6. Local sensitivity of system responses (T_{die} , PC/COP) to operating and design factors (q'' , Q_{liq} , v , S , L_e , L_c).	123
Figure 4.7. Response surface plots of the DOE-predicted die temperature (left) and power consumption (right) showing the combined effects of (a-b) liquid coolant and compressor speed, and (c-d) evaporator and condenser size.	124
Figure 4.8. P-h diagram comparing all demonstration cases (0, 1, 2) with $q'' = 86 \text{ W/cm}^2$ and $T_{die} = 55^\circ\text{C}$.	129
Figure 4.9. (a) Evaporator performance parameters (T_g , T_6 , ϵ) and (b) power consumption budget breakdowns for demonstration cases (0, 1, 2) for $q'' = 86 \text{ W/cm}^2$ and $T_{die} = 55^\circ\text{C}$.	130
Figure 4.10. (a) P-h diagram and (b) Power Consumption (PC) budget breakdowns comparing Cases 1 & 3.	133
Figure 5.1. Wireframe diagram of the miniature VCR assisted liquid-cooling system.	144
Figure 5.2. (a) Schematic of Validation Model configuration, (b) schematic of Heat Recovery Model configuration (note: the Heat Recovery Unit and Pump are shown for illustration purposes and are not modelled specifically here), and (c) P-h diagram schematic the Vapour-Compression thermodynamic cycle with description of the relevant operating points.	149
Figure 5.3. Schematic of the plate heat exchanger and key geometric specifications. This heat exchanger design is the evaporator in both the	151

Validation Model and Heat Recovery Model, as well as the condenser in the Heat Recovery Model.

Figure 5.4. Schematic representation and geometric specifications of the cross-flow air-cooled condenser used in the Validation Model. 152

Figure 5.5. Comparison of predicted and measured system (a) temperatures, (b) condenser and evaporator pressures and pressure ratio, and (c) Coefficient of Performance over a wide range of heat source powers, compressor speeds, air speeds, and liquid coolant flow rates. Predicted results obtained using the Validation Model. 156

Figure 5.6. Local sensitivity of system responses to operating and design factors. 159

Figure 5.7. Response surface plots for (a) T_{die} versus Q_{liq} & S , (b) T_{13} versus Q_{liq} & S , (c) PC versus Q_{liq} & S , (d) T_{die} versus L_e & L_c , (e) T_{13} versus L_e & L_c , and (f) PC versus L_e & L_c . For plots (a)-(c), q'' , L_e , L_c , Q_{hrf} & T_{12} , and for (d)-(f), the q'' , S , Q_{liq} , Q_{hrf} & T_{12} are fixed parameter settings. 160

Figure 5.8. P-h diagram comparing Cases 0 & 1, with $q'' = 86 \text{ W/cm}^2$, $T_{die} = T_{13} = 85^\circ\text{C}$. 166

Figure 5.9. Evaporator and condenser performance parameters (T_3 , T_6 , T_9 , T_{13} , ϵ) for demonstration Cases 0 and 1, for $q'' = 86 \text{ W/cm}^2$, $T_{die} = T_{13} = 85^\circ\text{C}$. 167

Figure 5.10. Power consumption budget breakdowns for demonstration Cases 0 and 1, for $q'' = 86 \text{ W/cm}^2$, $T_{die} = T_{13} = 85^\circ\text{C}$. 168

Figure 5.11. P-h diagram comparing Cases 1 & 2, with $T_{13} = 85^\circ\text{C}$. 170

Figure 5.12. Power consumption budget breakdowns for demonstration Cases 1 and 2, for $T_{13} = 85^\circ\text{C}$. 171

Figure 5.13. Heat pump Coefficient of Performance for demonstration 172
Cases 0, 1 and 2.

List of Tables

Table 1.1. Thesis Structure: Mapping of Research	18
Table 2.1. Uncertainties of measured quantities.	41
Table 2.2. Heat transfer uncertainties of FFAC and HLA systems.	42
Table 3.1. Uncertainties of measured quantities.	78
Table 3.2. Heat Transfer Uncertainties of VCR System	79
Table 4.1. Summary of design factors and response variables used in the DOE framework.	117
Table 4.2. Quantitative validation of the VCR-assisted liquid-cooling system mathematical model showing the Mean Relative Deviation (MRD) and Mean Absolute Relative Deviation (MARD) for each temperature measurement point.	120
Table 4.3. Comparison of experimentally derived and mathematically predicted optimum operating parameters from the DOE-based optimization for varying heat loads and target $T_{die} = 55^{\circ}\text{C}$ with ambient air at $T_{ambient} = 20^{\circ}\text{C}$.	121
Table 4.4. DOE optimization case study test matrix.	126
Table 4.5. DOE optimization case study results.	128
Table 5.1. Uncertainties of measured quantities.	147
Table 5.2. Heat transfer uncertainties of VCR system.	147

Table 5.3. Mean Absolute Relative Deviation (MARD) between the Validation Model predicted temperatures, pressures and system COP and experimentally determined values, over a wide range of operating conditions.	157
Table 5.4. Summary of design factors and response variables used in the DOE model.	158
Table 5.5. DOE optimization case study test matrix.	164
Table 5.6. DOE optimization case study results.	165

Nomenclature

A	Area	m^2
B_0	Boiling number	–
COP	Coefficient Of Performance	–
C_p	Specific heat capacity	$J/kg \cdot K$
dT/dx	Temperature Gradient	K/m
D	Diameter	m
F	Fractional Heat Exchanger Area	–
Fr	Froude number	–
G	Mass Flux	kg/m^2s
g	Acceleration due to gravity	m/s^2
h	Specific Enthalpy	kJ/kg
h_{air}	Air-side heat transfer coefficient	W/m^2K
h_{liq}	Liquid-side heat transfer coefficient	W/m^2K
h_{ref}	Refrigerant-side heat transfer coefficient	W/m^2K
I	Electrical Current	A
j	Colburn j-factor	–
k	Thermal Conductivity	W/mK
K	Friction factor constant	–
L	Length	m
LMTD	Log-Mean Temperature Difference	K
m	Friction factor exponent	–
$\dot{m}$	Mass Flow Rate	kg/s
Nu	Nusselt number	–
P	Pressure	Pa
PC	Power Consumption	W
P_{fan}	Fan Electrical Power	W
P_{comp}	Compressor Electrical Power	W
Pr	Prandtl Number	–

Q	Flow Rate	L/min
q''	Heat Flux	W/m ²
q	Heat Transfer Rate	W
R	Thermal Resistance	K/W
R'_T	Thermal Indicator	K/W
Re	Reynolds Number	—
T	Temperature	K, °C
UA	Overall Thermal Conductance	W/K
V	Voltage	V
v	Air Velocity	m/s
$\dot{V}$	Volumetric Flow Rate	m ³ /s
W_{pump}	Pump Hydraulic Work	W
We	Weber number	—

Greek Symbols

α	Louver angle	°
β	Chevron Angle	°
Δ	Difference	—
μ	Dynamic Viscosity	Pa·s
ρ	Density	kg/m ³
η	Efficiency	—
ξ	Electrical Efficiency	—
ε	Effectiveness	—
v	Specific Volume	m ³ /kg
σ	Surface Tension	N/m
χ	Vapour quality	—
δ	Uncertainty	—
ϕ	Two-phase multiplier	-

Subscripts

air	Air/ air-side
ambient	ambient
atm	atmospheric
c	Condenser
ch	Channel
clp	Cold Plate
comp	Compressor
CPU	Central Processing Unit
Cr	Critical
Cu	Copper
die	CPU Die
dis	displacement
e	Evaporator
eq	Equivalent
exp	Experimental
mod	Model
fan	Fan
FFAC	Fan-Fin Air Cooler
g	Vapor phase
HEX	Heat Exchanger
HLA	Hybrid Liquid-Air system
HP	Heat Pump
H	Hydraulic
hrf	Heat Recovery Fluid
IHS	Integrated Heat Spreader
in	Inlet
isen	Isentropic
LMTD	Log-Mean Temperature Difference
l	liquid phase

liq	Liquid/liquid-side
out	Outlet
pump	Fluid Pump
ref	Refrigerant/ refrigerant-side
sat	saturated
sc	sub-cooled
sh	superheated
th	Thermal
TIM	Thermal Interface Material
TP	Two-Phase
TTV	Thermal Test Vehicle
VCR	Vapour-Compression Refrigeration

Abbreviations

ACHE	Air-Cooled Heat Exchanger
AI	Artificial Intelligence
ANN	Artificial Neural Network
CMOS	Complementary Metal-Oxide-Semiconductor
COP	Coefficient of Performance
CPU	Central Processing Unit
DAC	Direct Air Capture
DAQ	Data Acquisition
DC	Direct Current
DIMM	Dual In-line Memory
dgt	Digit
DLC	Direct Liquid Cooling
DOE	Design of Experiments
EES	Engineering Equation Solver
FFAC	Fan-Fin Air Cooler
FPI	Fins Per Inch

FS	Full Scale
GPU	Graphics Processing Unit
GWP	Global Warming Potential
HFC	Hydrofluorocarbon
HG	Hold Value
HLA	Hybrid Liquid-Air
HP	Heat Pump
HPC	High Performance Computing
HRF	Heat Recovery Fluid
HRU	Heat Recovery Unit
HWC	Hot Water Cooling
IC	Integrated Circuit
IHS	Integrated Heat Spreader
LCC	Liquid Cooling Circuit
LMTD	Log-Mean Temperature Difference
MARD	Mean Absolute Relative Deviation
Max	Maximize
Min	Minimize
MPE	Multi-Port Extrusion
MRD	Mean Relative Deviation
MOO	Multi-Objective Optimisation
NTU	Number of Transfer Units
OAT	Organic Additive Technology
OFE	Oxygen-Free Electronic
PC	Power Consumption
PHE	Plate Heat Exchanger
PSU	Power Supply Unit
PUE	Power Usage Effectiveness
RPM	Revolutions Per Minute
RTD	Resistance Temperature Detector

SOO	Single-Objective Optimisation
TDP	Thermal Design Power
TEC	Thermoelectric Cooler
TG	Target Goal
TIM	Thermal Interface Material
TP	Two Phase
TTV	Thermal Test Vehicle
TXV	Thermostatic Expansion Valve
UC	Unconstrained
US	Ultrasonic
VCR	Vapour-Compression Refrigeration
VRM	Voltage Regulator Module

Chapter 1

Introduction

1.1. Background and Motivation

1.1.1. Data centre Energy Trends and Thermal Challenges

Data centres have become essential components of global digital infrastructure, supporting cloud platforms, high-performance computing, and increasingly complex data-driven applications. Their electricity demand has increased substantially in recent years. Across the EU-27, data centres consumed an estimated 45-65 TWh of electricity in 2022, representing approximately 1.8 - 2.6% of total EU electricity use [1]. More recent assessments indicate further growth, with some studies estimating that European data centres consumed around 96 TWh in 2024, or roughly 3% of the region's electricity [2]. As illustrated in Figure 1.1, forecasts for 2030 suggest a substantial rise in demand, driven by AI and hyperscale cloud expansion, with total data centre electricity consumption expected to nearly triple under current trajectories [3].

Because nearly all the electrical power consumed by computing hardware is dissipated as heat, cooling infrastructure must remove this thermal load continuously. The associated cooling energy demand is significant: mechanical and air-handling systems often account for 30 - 40% of a data centre's total electricity consumption, underscoring the strong coupling between thermal management and operational energy efficiency [4]. As compute density increases, cooling becomes one of the dominant contributors to overall facility power consumption and a major determinant of sustainability performance.

The root of this growing thermal burden lies in the evolution of modern compute hardware. For several decades, advances in processor performance were guided by two empirical scaling principles. Moore's Law, articulated in 1965, observed that the number of transistors on a chip doubles approximately every two years, enabling increased computational complexity and parallelism [5]. Dennard

scaling, introduced in 1974, showed that as transistors shrink, their operating voltage and current scale down proportionally, allowing switching energy to decrease and keeping power density roughly constant across successive generations [6]. The combination of these trends enabled rapid performance improvements without significantly increasing heat output, keeping cooling demands manageable.




Figure 1.1. Forecasted increase in European data centre electricity consumption and its share of total power demand between 2023 and 2030 [3].

However, as transistors approached nanoscale dimensions, the assumptions underpinning classical scaling began to fail. Supply voltage stopped scaling with feature size due to material and electrostatic limits, while leakage currents increased sharply. As a result, transistor switching energy no longer decreased as predicted, and power density began to rise with each new process node. This breakdown of Dennard scaling has been widely documented (Figure 1.2): despite continued increases in transistor density in line with Moore’s Law, the inability to reduce voltage led to higher chip power consumption and escalating thermal loads [7–10].

Modern CPU and accelerator architectures compensate for the loss of voltage scaling by increasing core counts, expanding execution width, and

employing aggressive boosting algorithms - all of which drive up total electrical power and heat generation. Server-class processors with sustained operating powers of 250-350 W [11–13], and specialised AI accelerators exceeding 400-700 W [12,14], are now common in European data centre deployments. When aggregated across densely packed racks, these devices create power densities far higher than those encountered even a decade ago.




Figure 1.2. Historical trends in microprocessor technology

Because nearly all this electrical energy appears as heat at the silicon junction, the cooling system must scale proportionally. Rising chip and rack power densities, combined with the plateauing of facility-level efficiency improvements (e.g., stagnating global average Power Usage Effectiveness of PUE~1.6 [15]), have created a situation in which traditional thermal management approaches are strained to their limits. Cooling capacity now directly constrains computational performance, energy efficiency, and long-term sustainability in European data centre environments.

These trends establish the broader motivation for examining the effectiveness of existing cooling methods. The next section considers air-based

cooling systems, the historical backbone of data centre thermal management, and outlines why they increasingly struggle to meet the demands imposed by modern high-density compute hardware.




Figure 1.3. Typical server node illustrating the principal heat-generating components within modern computing hardware, including the CPU, DIMMs, and Voltage Regulator Modules (VRMs) together with their relationship to server racks and data centre infrastructure.

A simplified representation of a modern server node and its principal heat-generating components is shown in Figure 1.3. The central processing unit (CPU), memory modules (DIMMs), and voltage regulator modules (VRMs) collectively contribute to the thermal load that must be removed by the cooling system. Multiple server nodes are typically installed within racks to form the computational infrastructure of modern data centres, where the combined heat generation of these components places increasing demands on cooling technologies.

1.1.2. Air-Cooling: Capabilities and Constraints

Air-cooling has historically been the dominant method of thermal management in data centre environments due to its simplicity, low cost, and compatibility with standard server architectures. Conventional systems rely on high velocity fans driving conditioned air through finned heat sinks mounted directly on processors. For earlier generations of CPUs with comparatively modest power levels, these configurations provided adequate heat removal while enabling dense rack layouts. As long as processor Thermal Design Power (TDP) levels remained below roughly 100-150 W, and rack power densities were modest (<10 kW), air-based cooling was able to maintain junction temperatures within acceptable limits [16,17].

However, the fundamental thermophysical properties of air impose physical limitations on its effectiveness as a cooling medium. Air's very low thermal conductivity (≈ 0.026 W/mK at 25°C), low specific heat capacity, and low density, severely constrain the achievable convective heat transfer coefficients, even under forced convection with high fan speeds [18,19]. As modern CPUs and accelerators began operating at 200-350 W and beyond, the thermal resistance between the silicon junction and the surrounding air required to maintain safe operating temperatures decreased significantly. Even with large finned heat sinks and aggressive airflow, maintaining safe junction temperatures became increasingly difficult, especially in dense server environments where inlet temperatures rise quickly under load [19].

In Europe, this challenge is compounded by rising rack power densities and limitations in facility cooling capacity. The European Commission's Joint Research

Centre notes that as hardware power levels rise, maintaining adequate airflow and low inlet temperatures becomes operationally expensive and energy intensive [20]. Higher air volumes demand greater fan power, which grows nonlinearly with pressure and airflow, resulting in diminishing returns and significant increases in energy consumption and acoustic output [21]. Efforts to extend the performance of air-cooling, such as increasing server fan speeds, using higher fin density heat sinks, or deploying advanced aisle containment strategies, provide only incremental improvements - they do not fundamentally overcome the main issue, which is air's poor thermal transport capability. Compounding this, studies show that even well designed air-cooled server racks experience inlet temperature non-uniformities, recirculation zones, and localised hotspots, all of which reduce effective thermal margins and can lead to thermal throttling under sustained high-power workloads [16]. As fan speeds escalate to compensate, noise constraints and power draw further limit practical operating ranges, while the cooling energy overhead grows disproportionately [22].

At the facility level, reliance on air as the primary heat transfer medium imposes substantial operational burdens. Moving sufficiently large volumes of chilled air through data halls becomes increasingly energy intensive as rack power densities surpass 20-30 kW per rack, and in some cases approach 40-60 kW per rack in high-performance environments [3,15]. Air-cooling, even when augmented with economisation or indirect evaporative cooling, struggles to maintain safe inlet temperatures under these conditions, particularly in urban European regions where climatic variations limit free-cooling windows [1].

These challenges underscore the inherent limitations of air-cooling in modern high density compute environments. Even with advanced airflow management, containment strategies, and increased fan power, air-cooling cannot overcome the fundamental constraints imposed by air's low thermal transport capacity. This conclusion is echoed in the ASHRAE Thermal Guidelines (4th Edition) [17], which document how airflow non-uniformities, rising inlet temperatures, and escalating rack power densities increasingly push air-based systems to the edge of their operational envelope. ASHRAE explicitly recognises that conventional air-cooling cannot adequately support future high-power IT equipment, and therefore supports

liquid-cooling as the necessary design pathway for higher power density server and rack deployments. As processors and accelerators continue to grow in power, the industry consensus reflected in ASHRAE guidance makes clear that liquid-cooling is not merely an alternative, but a necessary progression for sustaining reliable and energy efficient operation in next generation data centres.

1.1.3. Liquid-Cooling: Advantages and Emerging Limitations

As processor and accelerator power levels continue to rise, liquid-cooling has emerged as a practical and increasingly necessary alternative to traditional air-based thermal management. Liquids offer far superior thermophysical properties compared with air, including significantly higher thermal conductivity, density, and specific heat capacity [23–25]. These characteristics enable liquid-cooling systems to transport substantially more heat with smaller temperature differentials, lower pumping power, and reduced volumetric flow requirements. For high density compute environments, liquid-cooling provides a path to maintaining acceptable junction temperatures while limiting the energy overhead associated with coolant handling and transport.

Liquid-cooling can be implemented through several architectures, including rear-door heat exchangers, cold-plate direct-to-chip systems, and full immersion cooling [26–28]. Direct-to-chip liquid-cooling - now adopted among hyperscale operators and HPC installations - uses water-based or dielectric coolants circulated through cold plates mounted directly on the high-powered electronic components, like the CPU and GPU processors. By delivering the coolant closer to the heat sources, thermal resistances are reduced, and heat removal becomes more energy efficient, enabling facilities to support rack densities well above the limits of air-cooling [25,28,29]. Immersion cooling systems submerge the entire server assemblies in dielectric fluids, thus achieving cooling of the entire collection of electronic component within the server - from the lower powered DIMMs and VRM to the high powered CPUs and GPUs [24,27].

These system level benefits have led major technology providers and industry groups to position liquid-cooling as central to next generation data centre design. ASHRAE's Thermal Guidelines (4th Edition) formally introduces liquid-

cooling classes (W classes) that define acceptable coolant temperatures and operating envelopes for direct-chip and immersion architectures [17]. Market analyses from European and global operators similarly indicate that liquid-cooling is increasingly required to support high-density AI and HPC workloads, with many facilities already transitioning to hybrid liquid–air architectures to meet rising power demands [15].

Despite these advantages, liquid-cooling introduces new constraints and is not without limitations. The thermal resistance between the silicon die and the coolant still imposes a fundamental ceiling on achievable junction temperatures [25]. The liquid-side heat exchanger is just one component in a stack. Between the junction and the coolant, other components in the stack, including the semiconductor die, thermal interface material etc., contribute to the overall source-to-sink thermal resistance. Consequently, substantial improvements in cold plate design, such as enhanced microchannel geometries, jet impingement features, or low-conductivity materials, offer only marginal reductions in die-liquid thermal resistance. Moreover, the overall performance of direct-chip liquid-cooling is bounded by the coolant inlet temperature: even an optimised cold plate will struggle to maintain safe junction temperatures if the coolant enters at a relatively high temperature [30]. This emerging bottleneck has become increasingly apparent as processor power levels climb and heat fluxes approach the removal capacity of conventional water-based systems. While these challenges do not negate the thermal advantages of liquid-cooling, they highlight that liquid systems must be carefully engineered and integrated to be effective at scale.

As liquid-cooling becomes widespread, its limitations - particularly the constraints imposed by components within the source-to-sink stack, available supply of coolant, and coolant inlet temperatures - are becoming increasingly important [25]. These emerging challenges reveal that even liquid-cooling may struggle to support future processor generations unless the cooling medium itself is improved, or the coolant temperature is significantly reduced. This evolving landscape motivates deeper investigation into advanced liquid-cooling strategies and hybrid thermal management approaches capable of sustaining the demands of next generation compute hardware.

1.1.4. Coolant Temperature as a Limiting Factor

Liquid-cooling overcomes many of the limitations of air-cooling, but its performance is ultimately governed by the temperature of the coolant entering the cold plate. As discussed in above, a substantial portion of the junction-to-coolant thermal resistance originates within the processor package - primarily the die, TIM layers, and integrated heat spreader - and these elements are inherently difficult to modify. Because their design and material composition lie under the jurisdiction of the processor manufacturer rather than the cooling system integrator, they remain challenging to regulate or optimise at the system level. Consequently, once liquid cooled cold plates are effective to the point that package level thermal resistances dominate, the coolant inlet temperature becomes the most influential parameter determining the achievable junction temperature.

A further constraint arises from the heat rejection method used in typical liquid-cooling deployments. In most data centres, the liquid coolant loop ultimately rejects heat to ambient air via dry coolers, cooling towers or other air-based heat exchangers. This places a fundamental lower bound on attainable coolant temperature: the supply temperature cannot fall below ambient air temperature (plus a margin to account for heat exchanger effectiveness). Therefore, even with highly optimised cold plate performance, conventional liquid-cooling cannot deliver coolant below ambient conditions, meaning that its thermal capability is inherently limited by the environment in which the system operates.

In practical data centre deployments, coolant supply temperatures typically range from 20°C-35°C, a range chosen to maximise the use of regional free cooling and mitigate dependence on mechanical chillers [17]. While beneficial for facility level efficiency, these warmer coolant temperatures impose a narrow thermal margin for modern CPUs and accelerators, many of which exceed 300-600 W and operate at junction temperatures strongly influenced by coolant inlet conditions. Recent experimental studies reinforce this behaviour. Heydari *et al.* [26] showed that direct-chip cold plate systems cooling processors operating at realistically high workloads, provide limited junction temperature reductions when coolant temperatures are elevated, even when convective performance is optimised. Similarly, Li *et al.* [31] found that high heat flux components can exhibit unwanted

hotspots and temperature non-uniformity when supplied with warmer coolant, highlighting that package level conduction constraints amplify the sensitivity of junction temperature to coolant inlet temperature. A broad review of cold plate liquid-cooling technologies by Wu *et al.* [32] highlights the extensive research effort committed to channel geometry design, material selection, coolant formulation and flow distribution strategies in pursuit of improved thermal performance. These enhancements yield measurable reductions in cold plate thermal resistance; however, their impact is inherently bounded by the coolant temperature supplied to the cooling loop.

This leads to a central push-pull problem in contemporary thermal management of liquid cooled data centres: facility energy efficiency favours higher coolant temperatures, whereas device level reliability and performance benefit from colder coolant. As processor power densities and heat fluxes continue to increase, maintaining acceptable junction temperatures using elevated temperature liquid-cooling alone becomes increasingly challenging. Consequently, cooling architectures capable of selectively reducing coolant temperature - without imposing unfeasibly high energy penalty on the entire cooling plant - are receiving growing attention.

1.1.5. The Role of Active Cooling in Future Data Centres

The limitations imposed by elevated temperature liquid-cooling - particularly its dependence on ambient air for heat rejection - create situations where additional thermal margins cannot be achieved through passive means alone. When coolant supply temperature is bounded by outdoor conditions and heat rejection equipment (*i.e.* heat exchanger effectiveness), and when package level resistances limit the impact of cold plate performance, the only remaining engineering avenue is to actively reduce coolant temperature [30].

Active cooling methods, such as Thermoelectric Coolers (TECs) and Vapour-Compression Refrigeration (VCR), offer this capability by applying external work to lower coolant temperatures below ambient. TECs, while compact, are generally limited by low (sub-unity) COP and poor scalability for high-power applications [33]. In contrast, VCR systems provide substantially greater cooling capacity and far

higher energetic efficiency for the temperature ranges relevant to modern processors, making them a more practical choice for sub-ambient operation [30].

Previous investigations [34–39] have demonstrated that compact VCR systems can reliably achieve sub-ambient temperatures and support adequate cooling capacity for electronic devices, establishing their feasibility for CPU class thermal loads. Although earlier studies predominantly implemented direct-attach evaporators, they consistently showed that VCR technology is capable of delivering coolant temperatures significantly below ambient - offering a level of thermal control unattainable in conventional liquid-cooled systems. This ability positions VCR as a strong candidate for augmenting liquid-cooling loops in situations where additional, localised thermal margin is required.

However, this configuration of direct-attach evaporators presents several obstacles that limit its applicability to modern processors. When refrigerant evaporates directly on the heat source, regions of superheated vapour can develop within the evaporator passages, reducing the effective convective heat transfer coefficient and resulting in non-uniform cooling [40]. This is compounded by the relatively poor thermophysical properties of refrigerants, which restrict heat transfer effectiveness even under forced two-phase conditions [41]. Beyond these, direct attach architectures introduce practical integration challenges for cooling multiple heat sources, including piping complexity, leak protection and serviceability in dense server environments [17,42]. In addition, direct-attach evaporator VCR systems suffer from a fundamental scalability mismatch: the evaporators required for efficient two-phase boiling must have relatively large surface areas to keep heat fluxes within manageable limits, whereas modern CPUs and GPUs tend to concentrate very high heat fluxes onto comparatively small die and core regions. Designing an evaporator compact enough to fit within the processor's mechanical envelope, yet large enough to avoid excessive heat flux, dryout or flow instabilities, is challenging. This mismatch becomes even more severe for multi-die and chiplet architectures, where heat sources are both small and spatially separated. As a last point, characterizing two-phase flow behaviour is still a labour-intensive empirical process - at present it is far too complex to be accurately simulated for engineering design purposes. Considering the vast array of commercial processor design

architectures, together with the rate at which they evolve, this adds another layer of impracticability with regard to the use of direct attach evaporator cold plates in VCR-based cooling systems. Taken together, these underscore a clear research gap in current thermal management strategies and motivates the need for a cooling approach that can combine the performance benefits of liquid-cooling with the sub-ambient capability of VCR technology, while avoiding the limitations and impracticalities inherent to both.

1.1.6. Waste-Heat Recovery as an Emerging Driver for Advanced Cooling Architectures

Beyond the need to improve thermal performance, modern data centres present a growing opportunity for waste-heat recovery. Because nearly all electrical power consumed by CPUs, GPUs and associated electronics is converted to heat, data centres act as large, predictable thermal sources operating year-round. European sustainability frameworks - including the Energy Efficiency Directive (2023/1791) and the emerging Data Centre Sustainability Reporting Scheme [43–45] - explicitly encourage the utilisation of this waste heat wherever feasible, positioning heat recovery as an increasingly important design consideration for next generation facilities.

Traditional air-cooled data centres can only provide very low-grade waste heat. Return-air temperatures typically ranging from 30°C-45°C [46–48], offer limited reuse potential or require large heat pumps to lift temperatures to levels suitable for low-temperature district heating or industrial processes. The low density of air, combined with recirculation losses and temperature non-uniformities, further reduces heat quality, making meaningful recovery technically challenging and energetically costly [49,50].

Liquid-cooling can significantly improve the situation by consolidating waste heat into water-based loops with higher temperatures. This is possible when the cooling intensity provided by the cold plate heat exchanger is high enough that a process termed Hot Water Cooling (HWC) can be implemented. HWC refers to the scenario where ‘hot water,’ meaning coolant supply temperatures $\geq 40^{\circ}\text{C}$, can still maintain safe processor temperatures, thus producing heat recovery coolant outlet

temperatures $\geq 50^{\circ}\text{C}$ [51]. These temperatures are far more compatible with low-temperature district heating, building heating, and hot water reuse. Nevertheless, the use of data centre waste heat in district heating networks remains significantly constrained. These limitations arise not only from the requirement for close geographic co-location between the data centre and the heating infrastructure, but also from seasonal variability in heating demand and diurnal supply-demand mismatch. As a result, the aggregate annual utilisation of such waste heat is inherently limited. These constraints have driven growing interest in heat recovery solutions that operate on a continuous, year-round basis, offering the dual advantages of mitigating temporal supply-demand imbalance and enabling deployment independent of local district heating availability or site geography.




Figure 1.4. Simplified Direct Air Capture (DAC) process illustrating atmospheric CO₂ capture by a sorbent and subsequent thermally driven regeneration to release concentrated CO₂ for storage or utilisation.

A high potential round-the-year heat recovery end-use application for data centre waste heat is Direct Air Capture (DAC) [52]), a technology that extracts CO₂ directly from ambient air. In a typical DAC process, large volumes of air are drawn through a contactor containing a solid or liquid sorbent that selectively captures atmospheric CO₂ while allowing the remaining air to pass through. Once the sorbent becomes saturated, it must be regenerated to release the captured CO₂ and restore its capture capacity. This regeneration stage is typically achieved through the

application of thermal energy, producing a concentrated CO₂ stream suitable for storage or utilisation whilst allowing the sorbent to be reused in a continuous capture cycle [53]. A simplified schematic of the DAC process is shown in Figure 1.4.

Recent studies have identified data-centre waste heat as a potentially attractive energy source for DAC regeneration, creating opportunities to simultaneously improve data-centre sustainability and support carbon-removal technologies [52,54]. This, along with other potential applications, require supply temperatures in the 80°C-100°C range, which exceed what HWC liquid-cooling can provide without an additional temperature lift [53,54]. Since processor power densities are increasing whilst the safe operating temperatures are not, the temperature gap between *cooling needs* and *heat recovery requirements* is expected to widen. These motivate interest in engineering solutions capable of not only improving chip-level cooling performance, but also upgrading waste heat to higher temperatures to expand the options for potential reuse.

1.1.7. Towards Integrated Cooling and Heat-Recovery Systems

The preceding sections demonstrated that neither direct liquid-cooling nor direct-evaporator VCR systems may be sufficient to meet the combined thermal and operational demands of next generation high-power processors. Limitations relating to coolant temperature constraints, package-level thermal resistances, and the practical challenges of direct refrigeration at the chip, collectively indicate the need for a more scalable and capable cooling strategy. At the same time, existing heat recovery approaches in air and liquid-cooled systems are restricted to low or moderate grade temperatures, limiting their usefulness for high-value applications such as district heating or DAC regeneration.

Vapour-Compression Refrigeration (VCR) is uniquely positioned to overcome the key limitations identified in the previous sections. First, by placing the VCR evaporator in a secondary heat exchange loop, rather than directly on the chip, the system avoids the scalability and heat flux constraints inherent to direct-attach two-phase evaporators. In this solution, the primary loop uses a pumped water-based liquid coolant that can achieve much lower cold plate thermal resistances compared with boiling of refrigerants. The VCR evaporator then removes heat from

this primary loop whilst supplying sub-ambient coolant to the cold plate heat exchanger cooling the processor. This configuration therefore simultaneously resolves the coolant temperature problem of direct-attach liquid-cooling, and the discussed issues associated with direct-attach evaporators in VCR-based cooling systems. Second, the vapour-compression system inherently supports higher grade heat recovery. The condenser rejects not only the absorbed processor heat but also the compressor work, both of which appear as thermal energy at an elevated temperature due to the nature of the vapour-compression thermodynamic cycle. As a result, the vapour-compression system functions as a heat pump, upgrading low-grade heat to higher quality (energy and temperature) output, potentially suitable for DAC or similar applications. When incorporated into a hybrid liquid-cooling architecture, this has the potential to simultaneously deliver sub-ambient coolant to the cold plate and high-temperature heat at the condenser, addressing both advanced thermal management and sustainability objectives.

Hybrid vapour-compression assisted liquid-cooling therefore represents a promising pathway for next generation data centre thermal management. Its potential to improve both the thermal management and sustainability of data centres defines a clear and timely research opportunity, motivating the research undertaken in this thesis.

1.2. Research Objectives

The overarching objective of this research is to develop a novel miniature vapour-compression–assisted liquid-cooling technology suitable for integration within compact computer and server platforms for the thermal management of high-performance computing (HPC) processors. This work addresses the limitations of existing cooling technologies by investigating the potential of a hybrid vapour-compression assisted liquid-cooling architecture to deliver sub-ambient coolant temperatures, extend processor Thermal Design Power (TDP), and recover waste heat for reuse in emerging applications such as Direct Air Capture (DAC). The research combines experimental prototypes, system-level modelling, multi-objective optimisation and heat recovery analysis to establish an integrated design framework for advanced processor cooling. To achieve this primary research objective, the following specific sub-objectives are addressed:

1. To characterise existing CPU thermal management technologies, including air-cooling and hybrid liquid-air cooling, in order to (i) better understand their thermal and energetic performance, and (ii) establish baseline measurements on coolant temperature, thermal resistance and performance limits.
2. To design and develop a miniature Vapour-Compression - assisted liquid-cooling system, incorporating a primary single-phase coolant loop and a secondary VCR system capable of delivering sub-ambient temperature liquid coolant to a direct attach cold plate.
3. To experimentally characterise the capability of the VCR assisted liquid-cooling system, focussing on enhancing CPU thermal margins whilst maximizing system COP, with a key aim being to compare and contrast against conventional air and liquid-cooling technology.
4. To develop a mathematical model of the miniature VCR assisted liquid-cooling system, enabling detailed prediction of refrigerant thermodynamic states, component interactions and cycle performance, and to use this model to perform single and multi-objective optimisation of cooling capability and energy consumption.
5. To extend the developed mathematical modelling and optimization methodology to investigate the heat recovery potential of vapour-compression assisted liquid-cooling, evaluating the achievable temperature lift, recoverable heat output, and overall system energetics. A key aim is assessing the capability of the system to achieve heat recovery temperatures relevant for Direct Air Capture (DAC) applications.
6. To establish optimal operating strategies that balance cooling performance, energy efficiency and heat recovery potential, thereby determining practical operating envelopes for VCR-assisted liquid-cooling in data centre servers.
7. To establish a mathematical framework as a demonstration design tool for future miniature vapour-compression assisted liquid-cooling technological development, including the capacity for full system optimization.

1.3. Potential Research Impact

The expected outcomes of this research are underpinned by novel contributions to knowledge in the field of electronics cooling, data centre energy systems, and sustainable thermal management. The key contributions to knowledge are summarized below:

1. Deepening understanding of Vapour-Compression Refrigeration (VCR) technology in electronics cooling, particularly its integration within miniature sub-ambient liquid-cooling systems for high-power processors.
2. Development of new cooling system technology suitable for next generation electronic packages and system architectures, addressing the escalating thermal demands of HPC and AI platforms.
3. Creation of mathematical models and optimisation strategy that demonstrate a structured engineering methodology for the design and performance optimisation of sub-ambient electronics cooling systems, qualifying a new approach for evaluating and balancing cooling, energy consumption, and heat recovery performance.
4. Proving the concept that vapour-compression assisted liquid-cooling provides a feasible technological pathway for expanding the thermal, operational environment, and heat recovery envelopes for next generation computing platforms.

1.4. Thesis Structure

This thesis is structured in a “Thesis by Publications” format, in which each core chapter is written as a standalone research paper while contributing to a unified research narrative. The work follows a step-by-step progression from identifying the limitations of existing data centre CPU cooling technologies, through developing and characterising a miniature VCR-assisted liquid-cooling system, to advanced mathematical modelling and optimisation methods with waste heat recovery. Chapters 2 and 3 contain papers that have been published in international peer-reviewed journals, Chapter 4 is a paper currently under review, and Chapter 5 is

prepared for submission. In line with standard publication practice, each chapter follows the conventional structure of a journal article - including Abstract, Introduction, Methodology, Results, Discussion, and Conclusions - and each begins with a focused literature review establishing the specific knowledge gap addressed. Some repetition of background material is therefore unavoidable. Nevertheless, every chapter makes an independent and original contribution: Chapter 2 evaluates the performance and limitations of conventional cooling approaches; Chapter 3 experimentally investigates a miniature VCR-assisted liquid-cooling system; Chapter 4 develops and validates a mathematical VCR model and applies a structured optimization design methodology for cooling-focussed CPU thermal management; and Chapter 5 examines the feasibility of CPU cooling with DAC-suitable heat recovery. The final chapter synthesises the collective findings from these works and presents the overall conclusions of the thesis. A structured overview of the thesis is provided in Table 1.1.

Table 1.1. Thesis Structure: Mapping of Research

Chapter & Theme	Contribution	Key Insights
Chapter 1 Background, Research Gap & Objectives	Establishes advanced cooling needed for modern CPUs/GPUs	• Package thermal resistance dominates • Air & liquid-cooling limited by ambient temperatures • Conventional systems cannot deliver high temperature waste heat • Identifies dual research need: sub-ambient cooling with heat recovery
Chapter 2 Conventional Cooling Approaches (Experimental)	Evaluates existing air & liquid cooling performance	• Quantifies far superior thermal and energetic performance of single-phase liquid-cooling of CPUs over conventional air-cooling • Demonstrates ambient temperature constraints

		• Established that CPU cooling performance limitations of conventional air and liquid techniques • Confirms that sub-ambient cooling is high potential pathway for future systems
Chapter 3 VCR-Assisted Liquid-Cooling (Experimental)	Demonstrates feasibility of hybrid VCR-assisted liquid- cooling	• VCR enables sub-ambient coolant temperatures • Improves CPU thermal headroom • Reveals dependence on fluid-mover energetics; motivates optimisation
Chapter 4 VCR Mathematical Modelling & Optimisation (Numerical)	Builds and validates hybrid VCR assisted liquid-cooling mathematical model and optimization strategy; explores energy-cooling trade-offs	• Heat transfer – fluid dynamic - thermodynamic model developed and validated against experimental data • Model integrated in optimization software to predict optimal cooling - energy efficient operating strategies • Predicts low condenser temperatures, thus impractical for heat recovery applications • Low compressor pressure ratio and air-cooled condenser identified as limiting factors for heat recovery suitability
Chapter 5 Integrated Cooling, Energy & Heat-Recovery Optimisation (Numerical)	Examines integrated vapour-compression system for integrated cooling, energy & heat recovery	• Higher pressure-ratio compressor and refrigerant-to-liquid condenser heat exchanger implemented • Establishes feasibility of achieving CPU thermal targets with DAC suitable heat recovery temperatures (>80°C)

- Demonstrates mathematical framework for optimized system design
-

1.5. References

- [1] G. Kamiya, P. Bertoldi, Energy Consumption in Data Centres and Broadband Communication Networks in the EU, JRC Publications Repository (2024).
- [2] A. Larson, Data Centers Consume 3% of Energy in Europe: Understand Geographic Hotspots and How AI Is Reshaping Demand, POWER Magazine (2025).
- [3] Anna Granskog, Diego Hernandez Diaz, Jesse Noffsinger, Lorenzo Moavero Milanesi, and Pankaj Sachdeva, with Arjita Bhan and Sofia von Schantz, Unlocking the European AI revolution, McKinsey, October 2024.
- [4] Karthik Ramachandran, Duncan Stewart, Kate Hardin, Gillian Crossan, Ariane Bucaille, As generative AI asks for more power, data centers seek more reliable, cleaner energy solutions, Deloitte Insights, November, 2024
- [5] G.E. Moore, Cramming more components onto integrated circuits, Reprinted from Electronics, volume 38, number 8, April 19, 1965, pp.114 ff., IEEE Solid-State Circuits Society Newsletter 11 (2006) 33–35.
- [6] R. Dennard, F. Gaensslen, H.-N. Yu, V. Rideovt, E. Bassous, A. Leblanc, Design of Ion-Implanted MOSFET's with Very Small Physical Dimensions, Solid-State Circuits Newsletter, IEEE 12 (2007) 38–50.
- [7] M. Bohr, The Evolution of Scaling from the Homogeneous Era to the Heterogeneous Era, Technical Digest - International Electron Devices Meeting, IEDM (2011).

- [8] S. Borkar, Design challenges of technology scaling, IEEE Micro 19 (1999) 23–29.
- [9] S. Borkar, A.A. Chien, The future of microprocessors, Commun. ACM 54 (2011) 67–77.
- [10] M. Horowitz, 1.1 Computing’s energy problem (and what we can do about it), in: 2014 IEEE International Solid-State Circuits Conference Digest of Technical Papers (ISSCC), 2014: pp. 10–14.
- [11] Intel product specifications, <https://www.intel.com/content/www/us/en/ark.html>
- [12] Jonathon Evans, Michael Andersch, Vikram Sethi, Gonzalo Brito and Vishal Mehta, NVIDIA Grace Hopper Superchip Architecture Whitepaper, NVIDIA, November, 2022.
- [13] AMD EPYC™ 4th Gen 9004 & 8004 Series Server Processors – Details, AMD, December, 2025.
- [14] NVIDIA H200 GPU, NVIDIA, <https://www.nvidia.com/en-us/data-center/h200/> , December, 2025.
- [15] D. Donnellan, A. Lawrence, D. Bizo, P. Judge, J. O’Brien, J. Davis, M. Smolaks, J. Williams-George, R. Weinschenk, The Uptime Institute Global Data Center Survey, (2024).
- [16] R.R. Schmidt, E.E. Cruz, M. Iyengar, Challenges of data center thermal management, IBM Journal of Research and Development 49 (2005) 709–723.
- [17] D. Quirk, ASHRAE’s Data Center Thermal Guidelines— Air-Cooled Evolution, (2004).
- [18] T.L. Bergman, F.P. Incropera, A.S. Lavine, D.P. DeWitt, Introduction to Heat Transfer, John Wiley & Sons, 2011.
- [19] C. Patel, R. Sharma, C. Bash, M. Beitelmal, Thermal considerations in cooling large scale high compute density data centers, 2002.

- [20] Montevecchi, F., Stickler, T., Hintemann, R., Hinterholzer, S. (2020). Energy-efficient Cloud Computing Technologies and Policies for an Eco-friendly Cloud Market. Final Study Report. Vienna
- [21] R. and A.-C.E. American Society of Heating, High density data centers: case studies and best practices., American Society of Heating, Refrigerating and Air-conditioning Engineers, Atlanta, GA, 2008.
- [22] M. Beitelmal, D. Fabris, Servers and data centers energy performance metrics, *Energy and Buildings* 80 (2014) 562–569.
- [23] B. Ramakrishnan, C. Turner, H. Alissa, D. Trieu, F. Rivera, L. Melton, M. Rao, S. Chigullapalli, T. Getachew, V. Prodanovic, R. Lankston, C. Belady, V. Oruganti, Understanding the Impact of Data Center Liquid Cooling on Energy and Performance of Machine Learning and Artificial Intelligence Workloads, *J. Electron. Packag* 147 (2024).
- [24] D. Al Kez, A.M. Foley, F.W.B.M. Hasan Wong, A. Dolfi, G. Srinivasan, AI-driven cooling technologies for high-performance data centres: state-of-the-art review and future directions, *Sustainable Energy Technologies and Assessments* 82 (2025) 104511.
- [25] F.M. Naduvilakath-Mohammed, R. Jenkins, G. Byrne, A.J. Robinson, Closed loop liquid cooling of high-powered CPUs: A case study on cooling performance and energy optimization, *Case Studies in Thermal Engineering* 50 (2023) 103472.
- [26] A. Heydari, A.R. Gharaibeh, M. Tradat, Q. soud, Y. Manaserh, V. Radmard, B. Eslami, J. Rodriguez, B. Sammakia, Experimental evaluation of direct-to-chip cold plate liquid cooling for high-heat-density data centers, *Applied Thermal Engineering* 239 (2024) 122122.
- [27] P. Birbarah, T. Gebrael, T. Foulkes, A. Stillwell, A. Moore, R. Pilawa-Podgurski, N. Miljkovic, Water immersion cooling of high power density electronics, *International Journal of Heat and Mass Transfer* 147 (2020) 118918.

- [28] C. Roe, X. Feng, G. White, R. Li, H. Wang, X. Rui, C. Li, F. Zhang, V. Null, M. Parkes, Y. Patel, Y. Wang, H. Wang, M. Ouyang, G. Offer, B. Wu, Immersion cooling for lithium-ion batteries – A review, *Journal of Power Sources* 525 (2022) 231094.
- [29] R. Kong, H. Zhang, M. Tang, H. Zou, C. Tian, T. Ding, Enhancing data center cooling efficiency and ability: A comprehensive review of direct liquid cooling technologies, *Energy* 308 (2024) 132846.
- [30] F.M. Naduvilakath-Mohammed, M. Lebon, G. Byrne, A.J. Robinson, An experimental investigation of vapor compression refrigeration cooling and energy performance for CPU thermal management, *Case Studies in Thermal Engineering* 54 (2024) 103963.
- [31] X. Li, S. Guo, H. Sun, S. Yin, S. Liu, Z. Wang, Experimental study of the performance of liquid cooling tank used for single-phase immersion cooling data center, *Case Studies in Thermal Engineering* 63 (2024) 105386.
- [32] Z. Wu, G. Zhang, S. Lu, P. Leng, Y. Yu, J. Deng, W. Huang, A comprehensive review of cold plate liquid cooling technology for data centers, *Chemical Engineering Science* 310 (2025) 121525.
- [33] K. Liang, Z. Li, M. Chen, H. Jiang, Comparisons between heat pipe, thermoelectric system, and Vapour-Compression refrigeration system for electronics cooling, *Applied Thermal Engineering* 146 (2019) 260–267.
- [34] S. Trutassanawin, E.A. Groll, S.V. Garimella, L. Cremaschi, Experimental Investigation of a Miniature-Scale Refrigeration System for Electronics Cooling, *IEEE Transactions on Components and Packaging Technologies* 29 (2006) 678–687.
- [35] A.G.A. Nnanna, Application of refrigeration system in electronics cooling, *Applied Thermal Engineering* 26 (2006) 18–27.
- [36] Z. Wu, R. Du, Design and experimental study of a miniature vapor compression refrigeration system for electronics cooling, *Applied Thermal Engineering* 31 (2011) 385–390.

- [37] L.-Y. Jeng, T.-P. Teng, Performance evaluation of a hybrid cooling system for electronic chips, *Experimental Thermal and Fluid Science* 45 (2013) 155–162.
- [38] A. Poachaiyapoom, R. Leardkun, S. Wongwises, Miniature Vapor Compression Refrigeration System for Electronics Cooling, *Case Studies in Thermal Engineering* 13 (2018) 100365.
- [39] H. Zhou, J. Li, Development and analysis of a simple structured and economic miniature vapor compression refrigerator for cooling electronics in harsh environment, *Applied Thermal Engineering* 223 (2023) 120047.
- [40] S.G. Kandlikar, Heat Transfer Mechanisms During Flow Boiling in Microchannels, *J. Heat Transfer* 126 (2004) 8–16.
- [41] S. Garimella, A. Fleischer, J. Murthy, A. Keshavarzi, R. Prasher, C. Patel, S. Bhavnani, R. Venkatasubramanian, R. Mahajan, Y. Joshi, B. Sammakia, B. Myers, L. Chorosinski, M. Baelmans, P. Sathyamurthy, P. Raad, Thermal Challenges in Next generation Electronic Systems, *Components and Packaging Technologies*, *IEEE Transactions On* 31 (2009) 801–815.
- [42] M. Iyengar, M. David, P. Parida, V. Kamath, B. Kochuparambil, D. Graybill, M. Schultz, M. Gaynes, R. Simons, R. Schmidt, T. Chainer, Server liquid cooling with chiller-less data center design to enable significant energy savings, in: 2012 28th Annual IEEE Semiconductor Thermal Measurement and Management Symposium (SEMI-THERM), 2012: pp. 212–223.
- [43] Directive (EU) 2023/1791 of the European Parliament and of the Council of 13 September 2023 on energy efficiency and amending Regulation (EU) 2023/955 (recast) (Text with EEA relevance), 2023.
- [44] Data Centre Energy and Sustainability Performance Reporting Obligations, Department of Climate, Energy and the Environment, Ireland, February 2025.

- [45] Deirdre Nagle, Jay Sattin, Jessica Buttanhsaw, New Data Centre Sustainability Reporting Obligations Introduced, Mason Hayes Curran, August 2024.
- [46] H.M. Ljungqvist, M. Risberg, A. Toffolo, M. Vesterlund, A realistic view on heat reuse from direct free air-cooled data centres, *Energy Conversion and Management: X* 20 (2023) 100473.
- [47] K. Ebrahimi, G.F. Jones, A.S. Fleischer, A review of data center cooling technology, operating conditions and the corresponding low-grade waste heat recovery opportunities, *Renewable and Sustainable Energy Reviews* 31 (2014) 622–638.
- [48] M. Wahlroos, M. Pärssinen, J. Manner, S. Syri, Utilizing data center waste heat in district heating – Impacts on energy efficiency and prospects for low-temperature district heating networks, *Energy* 140 (2017) 1228–1238.
- [49] E. Oró, P. Taddeo, J. Salom, Waste heat recovery from urban air cooled data centres to increase energy efficiency of district heating networks, *Sustainable Cities and Society* 45 (2019) 522–542.
- [50] W.-X. Chu, C.-C. Wang, A review on airflow management in data centers, *Applied Energy* 240 (2019) 84–119.
- [51] P. Kasten, S. Zimmermann, M.K. Tiwari, B. Michel, D. Poulikakos, Hot water cooled heat sinks for efficient data center cooling: Towards electronic cooling with high exergetic utility, *Frontiers in Heat and Mass Transfer* (2010).
- [52] L. Zhang, Z. Zhao, B. Chen, M. Zhao, Y. Chen, Zero-Carbon Development in Data Centers Using Waste Heat Recovery Technology: A Systematic Review, *Sustainability* 17 (2025) 10101.
- [53] M. Fasihi, O. Efimova, C. Breyer, Techno-economic assessment of CO₂ direct air capture plants, *Journal of Cleaner Production* 224 (2019) 957–980.

- [54] C.D. Díaz-Marín, Z.J. Berquist, Flipping the switch: carbon-negative and water-positive data centers through waste heat utilization, *Energy Environ. Sci.* 18 (2025) 8403–8413.

Chapter 2

Closed Loop Liquid-Cooling of High-Powered CPUs: A Case Study on Cooling Performance and Energy Optimization

Abstract

This case study presents an experimental investigation of closed loop liquid-cooling of CPUs. The cooling system implements a pumped liquid supplying a cold plate fixed to the emulated CPU heat source and a remote Air-Cooled Heat Exchanger (ACHE). The Hybrid Liquid-Air (HLA) system employs 25% v/v Arteco-Freecor with distilled water solution as the primary coolant, and air as the secondary coolant. The thermal-hydraulic performance of the system is characterized over a range of operating conditions, including CPU heat loads ($12 \text{ W/cm}^2 - 72 \text{ W/cm}^2$), liquid volumetric flow rates ($0.5 \text{ L/min} - 8 \text{ L/min}$), and air velocities ($1 \text{ m/s} - 3 \text{ m/s}$). In order to provide a performance frame of reference, evaluations were also conducted with a conventional air-based fan-fin CPU cooling system. The results indicate that the HLA system significantly outperforms the traditional air-cooled fan-fin system, with much lower CPU temperatures and feasible power levels exceeding 300 W . Additionally, the study highlights the importance of the liquid flow rate and fan speed on the overall cooling performance. Finally, an optimal combination of liquid and air flow rates is identified for minimum energy consumption. Overall, these findings support the conclusion that conventional air-cooling is reaching its theoretical limit and closed loop liquid-cooling has the potential keep pace with the escalating power levels of modern CPUs. However, the additional energy cost of the cooling system should be carefully managed in order to maximize overall system Coefficient of Performance.

2.1. Introduction

The demand for high-performance computing has led CPU manufacturers to develop increasingly powerful processors. These CPUs continue to consume more power and generate higher heat fluxes, averaging about 70 W/cm^2 with Thermal Design Powers (TDPs) exceeding 150 W [1]. Intel's i9 7980XE, 56 Xeon Platinum 8470Q (formerly Sapphire Rapids), and 56 Core Xeon Platinum 9282, among others, exemplify this trend with power outputs of 280 W (overclocked), 350 W , and 400 W respectively [2]. Managing the heat generated by these CPUs has become a significant and important challenge, since effective cooling solutions are required to prevent thermal throttling or damage to the CPU.




Figure 2.1. Thermal resistance network of a CPU package.

A simplified thermal resistance network of a CPU packaging is shown in Figure 2.1. The core to coolant thermal resistance consists of five resistances in series, divided into two levels: level one includes the resistance R_{die} due to the conduction through the semiconductor die, the resistance R_{TIM1} across the Thermal Interface Material (TIM) between die and heat spreader, and the spreading resistance R_{IHS} within the Integrated Heat Spreader (IHS); level two consists of resistance R_{TIM2} across the TIM between IHS and base of the heat exchanger, and the resistance R_{HEX} associated with the heat exchanging technology to the coolant. Among these, the first-level resistance is challenging to regulate, as it falls under the jurisdiction of the chip manufacturer. Nevertheless, the resistances at the

second level, in particular R_{HEX} , can be engineered by electronic packagers to reduce R_{Total} .

The majority of CPUs, whether in PCs or servers, are directly cooled by fan-fin heat sinks. These heat sinks can have duct (parallel) or impinging (normal) air flow configurations [3]. Desktop PCs typically use impingement cooling, while servers use ducted cooling. As CPUs push the boundaries of performance and power consumption, the limitations of air-cooling become increasingly evident. The cooling threshold for air cooled fan-fin heat sinks can be as low as 40 W/cm^2 [4]. Despite efforts to optimize heat sink designs and enhance fan performance, there exists a critical point where air-cooling can no longer effectively cope with the escalating heat output of high-performance CPUs. This limitation can push power-hungry CPUs to their thermal limits, putting system performance, stability, and reliability at risk. The limitation of air-cooling becomes more pronounced when considering the non-uniform heat distribution of modern CPUs [5]. In addition to the heat transfer limitations of fan-fin air-cooling, other critical factors like reliability, leakage current, and harsh operating conditions (especially at elevated temperatures), have emerged as strong motivators for the adoption of liquid-cooling technology. Despite this, adoption rates have been slow. While air-cooling remains a practical and cost-effective solution for many everyday computing needs, it is important to recognize its limitations in the face of ever-evolving CPU thermal challenges. The pursuit of cutting-edge performance demands high performance thermal management technologies, and liquid-cooling is emerging as an obvious choice for modern high-powered CPUs.

Motivated by the superior thermophysical properties of liquids, single-phase liquid-cooling first emerged in the mid-1960s as a means to cool large mainframe computers. However, it was not until 1981 that Tuckerman and Pease [6] introduced the concept of direct-to-chip liquid-cooling, using microchannels to enable heat flux removal rates exceeding 790 W/cm^2 . However, with the advent and uptake of Complementary Metal Oxide-Semiconductor (CMOS) transistors in the 1990s, the computing industry shifted its attention away from liquid-cooling. This was primarily due to the reduced heat generated by CMOS technology, which made air cooling a feasible and cost-effective solution. Nonetheless, in recent years, there has been a

renewed interest in liquid-cooling technologies due to the substantial increase in power densities of transistors and the realization that air-cooling is approaching its theoretical limit.

Microchannel-based liquid cold plates have been extensively researched as a high potential cooling solutions for high heat flux electronics, primarily owing to the fact that smaller hydraulic diameter significantly increases the convective heat transfer coefficient, whilst the interconnected microchannels create high surface area density [7]. With a relatively slow start, there has been renewed interest in microchannel heat sinks for cooling of electronics since the early years of this century, and research efforts have centred around improving heat transfer capabilities while managing pressure drop penalties. Various concepts have been explored, including considerations of channel geometry, flow obstructions, wavy and zig-zag channels, and introducing periodic pinching and channel roughness [8–11]. This is not an exhaustive inventory of designs conceived to enhance the thermal-hydraulics of microchannel heat exchangers, yet serves as a testament to the availability of diverse design alternatives for engineering such technologies. Research highlights include the studies of Colgan *et al.* [12] who achieved high overall heat transfer coefficients $>200 \text{ kW/m}^2\text{K}$, and reported heat flux dissipation $>1000 \text{ W/cm}^2$ [13]. Nevertheless, the utilization of reduced channel sizes in microchannel heat sinks results in the trade-off of higher pressure drop and hydraulic power requirements [14]. Also, temperature non-uniformities can arise in microchannel heat exchangers due to the sensible heating of the fluid along the channel length, nor are they well-suited for managing hotspots [14]. The presence of temperature gradients in electronics has long been known as detrimental to their reliability due to the thermal stresses induced by differential thermal expansion [15].

Alternatively, liquid microjet array impingement is a complimentary approach to microchannels as they achieve remarkably high heat transfer rates without the need for surface area enhancements. For example, in Michna *et al.* [16] a microjet array dissipated heat flux of 1100 W/cm^2 , albeit from a very small target area. Similar to microchannel flows, the convective heat transfer coefficient increases as the size of the microjets decreases [17,18]. Despite being relatively less explored in the available literature, microjets are emerging as a formidable contender for high heat flux cooling applications due to their ability to generate substantially higher

convective heat transfer coefficients compared to traditional channel-based designs, and generally with lower pressure drop penalties [19]. The comparatively high convective heat transfer coefficients are attributed to the high velocity of the impinging jets, which results in a high stagnation heat transfer coefficient at the heated surface. By dispersing coolant across the entire heated surface, jet arrays also have the potential to greatly reduce undesired surface temperature gradients. Furthermore, microjet arrays can be designed to provide targeting of hotspots to improve surface temperature uniformity [19]. However, the heat transfer between neighbouring jets in the array tends to degrade rapidly, presenting an engineering challenge in achieving high cooling effectiveness over large surface areas at reasonable volumetric flow rates [20].

Hybrid microjet-microchannel systems have also been studied. Here, the liquid is discharged through the perforated orifice plate as multiple microjets into a microchannel feature on the cold plate, thus capitalising on the merits of each technology to achieve the high cooling intensity whilst mitigating high temperature gradients and hydraulic penalties. For example, Robinson *et al.* [21], achieved effective conductance levels $>400 \text{ kW/m}^2\text{K}$, at a moderate pressure drop ($< 100 \text{ kPa}$).

In a closed loop configuration, ensuring optimal performance relies not only on the cold plate heat sink but also on the integration of a reliable and compact pump to facilitate fluid circulation within the loop. Additionally, a liquid-air heat exchanger is required to dissipate the captured heat from the cold plate. Furthermore, improved heat transfer provided by the liquid cold plate comes at the price of the pressure drop and associated pumping power. This, and any additional fan power requirements, must be considered in the context of the overall energy consumption of the cooling system.

From a system integration perspective, existing research has primarily focused on the CPU alone, whereas the full source-to-sink thermal-hydraulics has been largely overlooked. Some studies aimed to address this gap by investigating the overall thermal resistance of liquid-cooling systems, including the remote heat exchanger and the interplay of all system components. A compact CPU cooling system was studied by Bintoro *et al.* [22], comprising a single water jet impingement

device and mini-channel remote heat exchanger, achieving an overall source-to-sink thermal resistance of approximately 0.35 K/W. Similarly, Liu *et al.* [23] and Chang *et al.* [24] used microjets and microchannel heat sinks respectively, coupled with a remote heat exchanger to achieve overall thermal resistances of about 0.17 K/W-0.23 K/W. Efforts by Whelan *et al.* [25] to design and develop a full closed loop CPU liquid-cooling system achieved an overall thermal resistance of 0.18 K/W, with a low pumping power of 1.5 W.

Considering the above discussion, the current study is motivated by the recognition that high-powered CPUs will necessitate cooling solutions that surpass the inevitable limitations of conventional fan-fin air-cooling. Single-phase liquid-cooling emerges as a very high potential option from technical and practical standpoints. However, existing literature highlights knowledge gaps concerning the system level assessment of closed loop liquid-based cooling technology. These gaps encompass full system level thermal-hydraulic performance characteristics across a range of operating conditions, as well as energy management considerations, particularly in terms of characterizing and optimizing the energy penalty associated with the full cooling system. Thus, the objective of this research is to conduct a comprehensive case study on CPU closed loop liquid-cooling, encompassing the thermal-hydraulics from the heat source to the heat sink. By performing a thorough parametric study across a sufficient range of system parameters, the research aims to fill gaps in the existing literature by gaining a deeper understanding of the parameter sensitivity, and how this influence overall performance of the end-to-end systems. Since this investigation compares the performance characteristics with traditional fan-fin technology, it also presents electronics packagers with missing information, as well as a methodology and characterization procedure, for making practical decisions when evaluating the limitations of conventional air-cooling. Finally, the investigation extends to optimizing the energy cost associated with closed loop liquid-cooling.

2.2. Experimental Apparatus and Data Reduction

2.2.1. Thermal Test Vehicle with Liquid-Cooled Cold Plate and Remote Air-Cooled Heat Exchanger




Figure 2.2. Wireframe diagram of the liquid-cooling system using remote air-cooled heat exchanger.

One key aim of this investigation is to evaluate the thermal-hydraulic performance characteristics of a pumped Hybrid Liquid-Air (HLA) system consisting of a cold plate and remote Air-Cooled Heat Exchanger (ACHE). To achieve this, an instrumented flow loop facility that emulates the practical implementation of this type of cooling solution was constructed. As Figure 2.2. indicates, the flow loop has three region of interest - the Thermal Test Vehicle (TTV), the primary (liquid) coolant loop,

and the ACHE. As will be discussed, the TTV is comprised of a variable heat source which is cooled by a liquid cold plate. The present study employs a coolant mixture comprising 25% v/v Artec-Freecor FTC® (an ethylene glycol-based Organic Additive Technology (OAT) heat transfer fluid containing corrosion inhibitors) [26] and 75% v/v distilled water as the primary coolant. This coolant mixture was selected to provide freeze protection and inhibit corrosion and microbial growth compared to water [26].

The primary loop includes a SAER CF gear pump with a AC10 series Parker Hannifin single-phase inverter drive that controls the pump speed to the desired liquid flow rate. The flow rate is measured by a non-invasive UF25B Cynergy ultrasonic flow meter with uncertainty of $\pm 3\%$ of the reading. Pressure drop measurements are taken across both the cold plate and the ACHE with a 26PCDFA6D Honeywell differential pressure transducer with uncertainty of $\pm 0.25\%$ FS. Similarly, temperature differentials are taken across both the cold plate and the ACHE with T-Type thermocouples, which were inserted into the flow via compression fittings. The temperature probes were calibrated against a platinum RTD reference probe to an uncertainty of a $\pm 0.2^\circ\text{C}$.

A Thermal Test Vehicle (TTV) was constructed to represent a CPU heat source. As depicted in Figure 2.3 (a), the TTV was machined from single block of ultra-pure Oxygen-Free Electronic (OFE) copper, with a lower neck region of $A_{\text{die}} = 20 \times 20 \text{ mm}^2$ square cross section, roughly the area of higher powered ($>150 \text{ W}$) commercial CPU dies (e.g. Intel i9 7980XE), which tend to be marginally larger than their lower-powered counterparts. Three OMEGALUX® CS series stainless steel sheath cartridge heaters of 6.35 mm diameter that could supply a maximum total power of 100 W each were inserted into the bottom of the 60 mm long lower neck section to provide heat to the upper 40 mm x 40 mm cross section and 2.5 mm thick overhang section, which represents a CPU Integrated Heat Spreader (IHS). The TTV section is then housed in a 3D printed case lined with insulation to minimized heat losses and provide a mounting fixture for the attached liquid-cooled cold plate.

The measurement of the heat flux from the lower neck to the IHS, as well as the temperature of the tip of the neck region, here representing the CPU die temperature, is facilitated by three 1.0 mm diameter T-type thermocouples spaced

at equal intervals along its length and inserted into drilled holes to a depth of 10 mm. As Figure 2.3 (b) illustrates, this allows the determination of the temperature gradient at the base of the IHS, dT/dx , such that the associated heat transfer to the IHS section, q_{CPU} , can be determined from,

$$q_{\text{CPU}} = -k_{\text{Cu}} A_{\text{die}} \frac{dT}{dx} \quad (3.1)$$

The temperature of the neck-IHS interface, T_4 (hereinafter T_{die}), represents the temperature of the CPU die and is determined from linear extrapolation,

$$T_{\text{die}} = T_3 - \left(\frac{q_{\text{CPU}}}{A_{\text{die}}} \right) \left(\frac{\Delta x_3}{k_{\text{Cu}}} \right) \quad (3.2)$$

The IHS above the heat source was cooled by a commercial cold plate (Nexalus Enflux™). As Figure 2.2 indicates, this cold plate consists of a plastic manifold that forces the liquid through an array of 0.7 mm orifices, creating an array of high velocity microjets that impinge on a nickel-plated copper baseplate. The principal geometric and hydraulic characteristics of the Nexalus Enflux™ cold plate are summarised in Appendix D. The cold plate employs a 49-jet impingement array designed to provide high local heat transfer coefficients while maintaining modest pumping power requirements. This cold plate was retained unchanged throughout the experimental and numerical investigations to ensure a consistent thermal interface, thereby allowing differences in performance to be attributed primarily to the investigated cooling architectures rather than variations in cold plate design. The cold plate was mounted to the TTV via four bolts tightened with a torque wrench to ensure correct and even pressure. A layer of Nexalus Hydronex™ thermal grease, a non-silicone thermal interface material (TIM) containing inorganic thermally conductive fillers and having a manufacturer-specified thermal conductivity of 4.1 W/m-K, is placed between the IHS and the cold plate to reduce contact resistance, typically referred to as TIM2.




Figure 2.3. (a) Schematic of the Thermal Test Vehicle (TTV) with mounted liquid-cooled cold plate, and (b) illustration of heated meter bar and the resulting linear Temperature vs. Distance plot.

With the heat transfer to the IHS measured, the overall die-to-coolant thermal resistance can be estimated as,

$$R_{th,clp} = \frac{\Delta T_{clp}}{q_{CPU}} \quad (3.3)$$

As suggested by Webb [27], the temperature difference across the cold plate, ΔT_{clp} , is calculated as the Log-Mean Temperature Difference (LMTD),

$$\Delta T_{clp} = \frac{T_{out} - T_{in}}{\ln \left(\frac{T_{die} - T_{in}}{T_{die} - T_{out}} \right)} \quad (3.4)$$

It must be noted that overall thermal resistance includes the resistance contributions of the IHS, TIM2 and the cold plate heat exchanger. In a commercial CPU package, additional thermal resistances associated with TIM1 needs to be included for accurate estimation of the die temperature.

Consistent with modern liquid-cooling technology currently deployed in high-performance desktop/workstation computing systems, the heat rejection to the air was facilitated via a remote air-liquid heat exchanger. In this investigation, the Nexalus NXQ 360 fan-cooled liquid-air flat tube heat exchanger of outer dimensions 400 mm x 130 mm x 60 mm, fitted with three 120 mm high static pressure fans, was used to transfer the heat from the liquid to the air. The heat exchanger is equipped with 16 flat copper tubes, each 2 mm wide, spanning its entire length. The copper fins of the radiator are non-louvered and have a density of 18 fins per inch (FPI). Thermocouples were positioned in the entrance and exit air streams in order to evaluate the temperature rise across the air-side of the heat exchanger. These measurements, along with those of the water-side allow for the calculation of the overall thermal resistance of the heat exchanger,

$$R_{th,ACHE} = \frac{1}{(UA)_{ACHE}} = \frac{\Delta T_{ACHE}}{q_{ACHE}} \quad (3.5)$$

where q_{ACHE} is the heat transfer rate from the cooling liquid to the air, which is the sum of the heat transfer from the TTV, q_{CPU} , and the hydraulic power to the coolant from the pump, W_{pump} , and ΔT_{ACHE} is the LMTD of the crossflow liquid-air heat exchanger and is calculated by,

$$\Delta T_{ACHE} = \frac{(T_{HEX,out} - T_{air,in}) - (T_{HEX,in} - T_{air,out})}{\ln \left(\frac{(T_{HEX,out} - T_{air,in})}{(T_{HEX,in} - T_{air,out})} \right)} \quad (3.6)$$

The fan speed was controlled by varying the supply voltage with a desktop DC power supply over the range of 6V - 14V. The penalty associated with the power

required for the liquid-side of the cooling system requires determination of the hydraulic pumping power across the cold plate and ACHE,

$$W_{\text{pump}} = \dot{V}_{\text{liq}}(\Delta P_{\text{liq,clp}} + \Delta P_{\text{liq,ACHE}}) \quad (3.7)$$

where $\dot{V}_{\text{liq}}$ is the liquid volumetric flow rate and ΔP_{liq} is the associated pressure drop. It is important to note that the pressure drops arising from the valves, T-joints, flow meters are not considered here as these would be absent in a real-life application. The total pumping power along with the total electrical fan power,

$$P_{\text{fan}} = IV \quad (3.8)$$

allows for the estimation of the overall system Coefficient of Performance (COP),

$$\text{COP}_{\text{ACHE}} = \frac{q_{\text{CPU}} + W_{\text{pump}}}{W_{\text{pump}} + P_{\text{fan}}} \quad (3.9)$$

It should be noted that the coefficient of performance (COP) adopted in this work is intended as a system-level energy performance indicator rather than a classical thermodynamic COP. The parameter is defined as the ratio of the thermal load managed by the cooling system to the total electrical power required for its operation and was selected to facilitate comparison between fundamentally different cooling architectures, including air-cooling, liquid-cooling, and vapour-compression-assisted cooling systems. By accounting for both cooling capacity and auxiliary power consumption, the metric provides a convenient measure of the energetic effectiveness of the cooling system. However, the COP values reported herein are specific to the definition adopted in this study and may not be directly comparable with COP values reported elsewhere in the direct-chip cooling literature, where alternative definitions and performance metrics are frequently employed. Consequently, COP should be interpreted alongside thermal resistance, die temperature, and maximum sustainable heat load when assessing overall cooling performance.

2.2.2. Fan-Fin CPU Cooler

In order to establish a performance frame of reference for comparison with the liquid-cooled system described above, commercial off-the-shelf Fan-Fin Air-Cooler (FFAC) (ARCTIC Alpine 11 Intel CPU Cooler) was also characterized. The FFAC used was typical of most commercial impingement type heat sinks, with aluminium fins measuring 95 mm x 98 mm x 70 mm and equipped with a 92 mm fan. The heat sink was affixed to the Integrated Heat Spreader (IHS) of the TTV via mechanical fastening, with the same layer of thermal grease (TIM2) placed between the IHS and the aluminium heatsink base, as depicted in Figure 2.4. The heat sink had a total of 35 fins, each of 0.6 mm thickness and 34.5 mm height, with 1.8 mm spacing between them.




Figure 2.4. Wireframe diagram of the impingement flow Fan-Fin Air-Cooler (FFAC) system.

The FFAC selected for this research was intended to provide a realistic and commercially representative benchmark for conventional direct-attach air-cooling rather than the absolute performance limit of air-cooling technology. The selected ARTIC Alpine 11 Intel CPU cooler represents a class of compact fan-fin coolers commonly employed in desktop, entry-level servers, and other cost-sensitive computing platforms. As such, it provides a practical baseline against which the performance of the Hybrid Liquid-Air (HLA) cooling system can be assessed, with further comparison to the Vapour-Compression Refrigeration (VCR) cooling system presented in subsequent chapters.

It is acknowledged that larger heat sinks, alternative fin geometries, heat pipes, or vapour chamber technologies may further reduce thermal resistance. However, the resulting performance improvements are typically accompanied by engineering trade-offs including increased size, weight, cost, airflow requirements, pressure drop, fan power consumption, and acoustic noise. Furthermore, as heat sink size increases, progressively smaller reductions in thermal resistance are generally achieved because the air-side convective resistance becomes increasingly dominant. Consequently, while higher-performance air-cooling solutions are possible, they remain fundamentally constrained by the thermophysical limitations of air and by the requirement to reject heat directly to the ambient environment.

Therefore, the objective was not to compare against an optimised air-cooling solution, but rather to establish a realistic baseline representative of conventional air-cooling architectures and to quantify the performance gains achievable through liquid-cooling and VCR-assisted cooling technologies.

In characterizing the thermal performance of the FFAC, the source-to-sink thermal resistance ($R_{th,FFAC}$) is calculated as,

$$R_{th,FFAC} = \frac{\Delta T_{FFAC}}{q_{CPU}} \quad (3.10)$$

where ΔT_{FFAC} is the LMTD of the cold plate heat sink and is calculated by

$$\Delta T_{\text{FFAC}} = \frac{T_{\text{air,out}} - T_{\text{air,in}}}{\ln \left(\frac{T_{\text{die}} - T_{\text{air,in}}}{T_{\text{die}} - T_{\text{air,out}}} \right)} \quad (3.11)$$

By utilizing a benchtop DC power supply, the fan speed was regulated, and the total power was computed similar to Eq. (3.8), such that the system's Coefficient of Performance (COP) can be determined from,

$$\text{COP}_{\text{FFAC}} = \frac{q_{\text{CPU}}}{P_{\text{fan}}} \quad (3.12)$$

2.2.3. Uncertainty Analysis

The uncertainties associated with the primary measurement instruments are provided in Table 2.1. To determine the propagated uncertainty of calculated quantities, standard methodologies outlined by Taylor [28] were followed, assuming that the uncertainties were random and independent. In cases where a calculated quantity Z relied on measured quantities X , Y , and possibly others, the uncertainty δZ was determined using the following prescribed procedures.

$$\delta Z = |Z| \sqrt{\left(\frac{\delta X}{|X|} \right)^2 + \left(\frac{\delta Y}{|Y|} \right)^2 + \dots} \quad (3.13)$$

Table 2.1. Uncertainties of measured quantities.

Measurement	Instrument	Abs. Uncertainty
Temperature, T	T-type Thermocouple	$\pm 0.2^\circ\text{C}$
Liquid Flow Rate, $\dot{V}_{\text{liq}}$	Ultrasonic Flow Meter	$\pm 3\%$ reading
Pressure Drop, ΔP_{liq}	Diff. Pressure Sensor	$\pm 0.25\%$ FS
Air Velocity, v	Vane Anemometer	$\pm 2\%$ reading
Voltage, V	Digital Multimeter	$\pm (0.09\% + 2\text{dgt})$
Current, I	Digital Multimeter	$\pm (1.0\% + 3\text{dgt})$

Table 2.2. Heat transfer uncertainties of FFAC and HLA systems.

Heat Transfers	Uncertainty (%)			
	FFAC		HLA	
	Min.	Max.	Min.	Max.
T_{die} ($^{\circ}\text{C}$)	± 0.17	± 0.39	± 0.35	± 0.59
q_{CPU} (W)	± 2.0	± 3.7	± 2.0	± 3.6
P_{fan} (W)	± 3.3	± 5.8	± 2.3	± 4.0
W_{pump} (W)	-	-	± 1.4	± 1.8
COP (-)	± 4.7	± 8.1	± 3.8	± 6.2

The uncertainty associated with the rate of heat transfer through the heated meter bar system requires some special attention, as it depends not only multiple temperature and spatial errors, but also the regression fitting to these to determine dT/dx . Similarly, this error propagates to the estimation of T_{die} , which must also be quantified. To account for this, the Monte-Carlo calculation method first described in Kempers *et al.* [29] and later detailed in Kempers and Robinson [30] was implemented. Table 2.2 gives the propagated uncertainty for all derived quantities, which follow the standard methodology.

2.3. Results and Discussion

2.3.1. Thermal Characterisation

2.3.1.1. Fan-Fin Air-Cooling System

Figure 2.5 shows that, for the FFAC system, as the heat flux supplied from the TTV increases from 12.25 W/cm^2 to 72 W/cm^2 the net thermal resistance remains relatively constant within their absolute uncertainties. Increasing the air velocity across the heat sink lowers its thermal resistance from about 0.32 K/W to 0.27 K/W , indicating marginally improved heat transfer. Figure 2.6 shows the expected cooling performance of the FFAC system under various heat loads for an ambient temperature of 35°C , the standard room temperature recommended by

Intel [31]. Under extreme loads, the average die temperatures can reach as high as 80°C-85°C, and this is regarded as a 'safe' operating temperature limit [32]. The throttling temperature limit may differ depending on the CPU type, typically falling within the range of 100-105°C [33]. In this study, the safe and throttling temperature limits of 83°C and 102°C respectively, are chosen when comparing cooling technologies. To predict the die temperature closer to what would be expected in a CPU package, an approximation for the TIM1 thermal resistance, using the $R_{TIM1} = 0.1 \text{ K-cm}^2/\text{W}$, was included [4].




Figure 2.5. Thermal resistance of FFAC system with varied heat loads; $T_{ambient} = 35^\circ\text{C}$.




Figure 2.6. Expected cooling performance of a FFAC at maximum fan Speed with varying heat loads: $T_{\text{ambient}} = 35^{\circ}\text{C}$ and $R_{\text{TIM1}} = 0.1 \text{ K-cm}^2/\text{W}$

For a normal fan speed of 1.8 m/s, the safe die temperature limit of 83°C was breached at approximately 35 W/cm^2 , which is close to that reported in the literature [4]. The temperature at the maximum heat flux studied (72 W/cm^2) was at least 53 K higher than the safe limit, and even well above the throttling temperature of 102°C [33]. Consistent with the thermal resistance behaviour, there is a noticeable yet marginal effect of air velocity on the die temperature which diminishes with increased velocity. Considering the diminishing gain in the heat transfer, together with the escalation in fan noise and power consumption, it is not worthwhile to increase the air speed beyond approximately 1.8 m/s.




Figure 2.7. Power Consumption (PC) & Coefficient of Performance (COP) versus heat load for FFAC system.

As Figure 2.7 indicates, increasing the heat flux results in an increase in the Coefficient of Performance (COP) of the system, ranging approximately from ~15 to 90 for settings that achieve lower than the throttling temperature. This result is expected since, despite the increase in heat transfer rate, the fan power required to cool CPU was kept constant. However, as Figure 2.7 shows, higher fan speeds lead to a significant decrease in the COP. This occurs because, even though the heat flux remains constant, the fan power consumption nearly doubles when the air velocity changes from 1.3 m/s to 2.1 m/s, resulting in the COP being halved. Upon closer examination of the temperature and COP plots, it was observed that surpassing an air velocity of 1.8 m/s does not provide substantial cooling benefits. Instead, it leads to increased power consumption, lower COP values, and heightens fan noise. This highlights a key point which has not been addressed sufficiently in the literature: a balance needs to be struck between cooling requirements and the

associated penalties, like energy consumption and noise. Each aspect must be considered holistically in the development of a cooling system.

2.3.1.2. Hybrid Liquid-Air (HLA) Cooling System




Figure 2.8. Thermal resistance of HLA cooling system with varied heat loads; $T_{ambient} = 35^{\circ}\text{C}$.

Figure 2.8 shows that for the selected nominal settings (max fan speed & $Q_{liq} = 4 \text{ L/min}$), the thermal resistance of the cold plate and ACHE remains almost constant, as the heat flux increases from 12.25 W/cm^2 to 72 W/cm^2 . It is worth noting that the thermal resistance of the cold plate is roughly ten times that of the ACHE, which supports the focus on cold plate research in the literature. Notably, the net thermal resistance of the liquid-cooling system is about 0.11 K/W , which three times lower than the FFAC, which is significant.




Figure 2.9. Expected cooling performance of HLA cooling system with varying heat loads for a nominal flow rate of $Q_{liq} = 4$ L/min and fan speed of $v = 3$ m/s; $T_{ambient} = 35^\circ\text{C}$ and $R_{TIM1} = 0.1$ K-Cm²/W.

The expected die temperature with increasing heat flux are plotted in Figure 2.9, again considering $T_{ambient} = 35^\circ\text{C}$ and approximate TIM1 thermal resistance. As the primary and secondary coolant temperatures experience only slight increases over the range of considered heat fluxes, T_{die} increases linearly with the heat transfer rate, attributable to the constant source-to-sink thermal resistance. The HLA system maintains temperatures well within the safe operating limits for the full range of high heat fluxes tested, and throttling is not expected until power levels of close to 450 W. When comparing the performance of the HLA and FFAC systems, it was found that the HLA system is capable of cooling the die by at least 60K more effectively than the fan-based system for the highest tested heat flux of 72 W/cm².




Figure 2.10. Pressure drop curves of the HLA system for varying heat loads.

Figure 2.10 plots the HLA system pressure drops at varying heat load for the nominal flow rate of 4 L/min across both the cold plate and ACHE. The cold plate demonstrates an expectedly higher pressure drop compared to ACHE (approximately 16 kPa compared to 2 kPa), though these combine to a pumping power of only 1.1 W. It is worth noting that fan power at the ACHE for this maximum air velocity is nearly six times greater than the hydraulic power necessary to circulate the liquid in the loop, as shown in Figure 2.11. Therefore, previous research efforts that focused on enhancing cold plate heat transfer capabilities whilst managing the pressure drop is justified.

As seen in Figure 2.11, the COP of the system is observed to rise linearly, as expected, from approximately 5 to 36 over the range of TTV powers tested. A slight decrease in hydraulic power is observed due to the lower pressure drop in the liquid loop caused by a slightly lower coolant viscosity at the higher operating

temperatures, as seen in Figure 2.10. Compared to the FFAC, the COP of the HLA system is approximately 2.5X lower for settings that achieve lower than the throttling temperature. However, this penalty must be considered in the context of the benefits associated with the lower operating CPU temperatures and extended power range afforded by the HLA system.




Figure 2.11. Power Consumption & Coefficient of Performance profile of the HLA system.

2.3.2. Liquid Flow Characterisation - HLA System

Figure 2.12 shows the variation of thermal resistance with liquid flow rate for a fixed maximum heat flux and fan speed. The plot shows that as the liquid flow rate increased from 0.5 L/min to 8 L/min, the thermal resistance decreased, though its effect is much pronounced for the cold plate compared to the ACHE. However,

beyond around 4 L/min, the flow rate does not have significant effect on the thermal resistance. This is also noticeable from the temperature plots in Figure 2.13, whereby any increase in coolant flow rate beyond 4 L/min has diminishing influence on the die temperature, which are all predicted to be below the safe operating temperature for this maximum heat flux rest case. This diminishing return with liquid flow rate is expected since (i) jet array impingement heat transfer increases as $Nu \propto Re_d^n$, with $n \sim 0.5$ [17], and (ii) as the convective resistance decreases, those of the IHS and TIM2 become increasingly dominant.




Figure 2.12. Thermal resistance of HLA system for varying liquid flow rates measured at constant heat flux and air velocity; $T_{ambient} = 35^{\circ}C$.




Figure 2.13. Expected cooling performance curves of HLA system for varying liquid flow rates at constant heat flux and air velocity; $T_{\text{ambient}} = 35^{\circ}\text{C}$ and $R_{\text{TIM1}} = 0.1 \text{ K-cm}^2/\text{W}$.

However, as shown in Figure 2.14, the pressure drop across the cold plate and ACHE increases with the expected parabolic profile, with the former dominating the total system pressure drop and thus hydraulic power. As such, the diminishing gains in heat transfer associated with flow rates beyond about 4 L/min has associated with it increasingly severe hydraulic penalties. Again, this illustrates that, when considered as a full system, careful design consideration must be made to ensure adequate performance with managed hydraulic penalties. This is particularly important considering the impact on the overall energy dynamics of the HLA system. This is displayed in Figure 2.14, where increasing flow rates from 4 L/min to 8 L/min leads to about a 3.5-fold increase in the hydraulic power, which negatively affects the system COP, as shown in Figure 2.15, with little improvement in the cooling performance (Figure 2.13). However, this is moderated to some degree by the fan

power (~ 6.6 W), which is notably higher than the hydraulic power for liquid flow rates up to 4 L/min. This causes the COP to level-off at lower liquid flow rates, and to a lower value for decreasing heat input levels (see Figure 2.15).




Figure 2.14. Pressure drop of HLA system for varying liquid flow rates at constant heat load.




Figure 2.15. Coefficient of Performance of HLA system for three different heat fluxes with varying liquid flow rates and constant air velocity.

2.3.3. Fluid Characterisation - Air Side

Keeping the heat flux constant (72 W/cm^2) and liquid flow rate at its nominal 4 L/min , the influence of the air velocity (fan speed) was varied between 1 to 3 m/s to investigate its overall influence on system performance. As evident from Figure 2.16, increasing the air velocity has much less significant influence on the ACHE thermal resistance compared with the effect the liquid flow rate has on the cold plate resistance (Figure 2.12). It is thus unsurprising that the predicted die temperature is nearly constant with increasing fan speed, as seen in Figure 2.17. The same cannot be said about the power consumption and related COP, which depicted in Figure 2.18, shows the system COP decreasing steeply with increased fan speed. This is due to the increased power consumption of the fan, which increases about 5.5-fold. Notably, Figure 2.18 shows that the fan speed can be almost as low as the pumping

power at the lowest fan speed setting, and this occurs with only a minor drop in the die temperature (Figure 2.17).




Figure 2.16. Thermal resistance of HLA system for varying air velocities at constant heat flux and liquid flow rate; $T_{ambient} = 35^{\circ}\text{C}$.

This investigation of closed loop liquid-cooling on a high-powered CPUs showed significantly higher cooling capacity than the traditional FFAC system. For a given setting of the fans and pump, the overall thermal resistance is relatively insensitive to the applied heat flux. Increasing the liquid flow rate improved the heat transfer to a point, after which there were diminishing returns in terms of higher power consumption. Lastly, increasing the fan speed had minimal impact on the overall heat transfer while significantly increasing power consumption and decreasing the COP. These findings highlight the importance of effective management of the fluid movers to optimize energy dynamics in liquid-cooling systems. Optimising power consumption of the HLA system is crucial, as it remains the main drawback

compared to conventional FFAC systems, highlighting the need for a comprehensive power optimization study.




Figure 2.17. Temperature profile of the HLA system for varying air velocities measured at constant heat flux and liquid flow rate; $T_{\text{ambient}} = 35^\circ\text{C}$ and $R_{\text{TIM1}} = 0.1 \text{ K-cm}^2/\text{W}$.




Figure 2.18. Power Consumption and Coefficient of Performance of the HLA System for varying air velocities at constant heat flux and liquid flow rate.

2.3.4. Cooling Energy Optimization

To illustrate the potential for optimizing the cooling energy requirement of the HLA system, cases are considered where the fluid movers are controlled in such a way to maintain a constant die temperature. Practically, CPUs operate from idle to their TDP, and maintaining a constant die temperature would be beneficial since it would reduce thermal cycling over large computational load swings. Two case studies are considered, one for maintaining $T_{\text{die}} = 75^{\circ}\text{C}$ and a more aggressive case where $T_{\text{die}} = 55^{\circ}\text{C}$, each for the more common situation where $T_{\text{ambient}} = 20^{\circ}\text{C}$. The full heat flux range from 6 W/cm^2 to 72 W/cm^2 is then considered for the primary liquid coolant flow rates between 0.5 L/min - 8 L/min and the secondary air coolant velocities of 1 m/s - 3 m/s .




Figure 2.19. HLA Power Consumption (PC) surface plots for maintaining the die temperature at (a) $T_{die} = 75^{\circ}\text{C}$ and (b) $T_{die} = 55^{\circ}\text{C}$, by varying liquid flow rate and air velocity.

Figure 2.19 (a) & (b), plot the total cooling power consumption to maintain die temperatures of 75°C and 55°C respectively, over the range of die heat fluxes and as a function of the pump and total fan flow required to sustain the setpoint temperature. To effectively maintain the constant die temperature at a given heat flux, Figure 2.19 shows that the minimum cooling power consumption adapts to a higher liquid flow rate and lower air velocity (fan speed). For example, considering the 75°C / 60 W/cm² case, the same overall cooling performance can be achieved with a liquid flow rate of 0.52 L/min and air velocity of ~3.0 m/s, as it can with a higher liquid flow rate of 2.2 L/min and lower air velocity of 1.0 m/s. Here the total cooling power drops from 5.2 W to 1.8 W, which is better than what can be achieved by the FFAC system, despite the fact that the FFAC system cannot cool sufficiently at this heat flux level (Figure 2.7). Referring to the COP surface plot in Figure 2.20. (a), the resulting gain for this case is from 52.1 to 156.4, which is significant.

For the more aggressive cooling case of 55°C / 60 W/cm², Figure 2.19 (b) and Figure 2.20 (b), it is optimal to operate the HLA system with the liquid and air fluid movers set at 4 L/min and 1 m/s respectively, resulting in an energy consumption of 3.3 W (COP~75). In contrast, utilizing the maximum air speed of 3 m/s, a liquid flow rate of only 1.9 L/min is required, though the energy consumption is 2 times higher, with a commensurate drop in COP. Therefore, opting for lower fan speed requires higher liquid flow rates to deliver equivalent cooling performance, though consuming comparatively less power, which improves the COP. It is interesting to note that the maximum air-speed of 3 m/s is never required for the two setpoint test cases, and apart from the highest heat flux, only ~ 1 m/s (lowest fan speed setting) is generally required for minimum cooling power and maximum COP. Thus, the temperature can be regulated, at minimum cooling power expense, by controlling the liquid flow rate. In doing so, high COPs can be achieved for the range of CPU heat fluxes between 24 W/cm² to 72 W/cm², ranging between COP ~ 97-150 for the $T_{die} = 75^{\circ}\text{C}$ setpoint, and COP ~ 60-73 for the $T_{die} = 55^{\circ}\text{C}$ setpoint. Importantly, this is comparable to a FFAC cooling system, whilst significantly expanding the maximum TDP capability of CPUs.




Figure 2.20. HLA Coefficient of Performance (COP) surface plots for maintaining the die temperature at (a) $T_{die} = 75^\circ\text{C}$ and (b) $T_{die} = 55^\circ\text{C}$, by varying liquid flow rate and air velocity.

2.4. Conclusions

This case study aimed to investigate the cooling performance of a closed loop liquid-cooled system for advanced thermal management of high-performance CPUs. For an emulated $2 \times 2 \text{ cm}^2$ CPU die and range of input power and heat fluxes up to 288 W and 72 W/cm^2 respectively, the experimental investigation measured key cooling performance metrics of a Hybrid Liquid Air (HLA) system, such as die temperature and thermal resistance, over ranges of liquid flow rates (0.5 L/min - 8 L/min) and fan speeds (air velocities of 1 m/s to 3 m/s). The cooling pressure drop, and power penalty was also measured. To provide a comparison, a traditional Fan-Fin Air-Cooled (FFAC) system was also tested, and the following conclusions can be drawn;

- (1) The HLA system outperforms the FFAC system in terms of source-to-sink thermal resistance and associated lower CPU temperatures.
- (2) The feasibility limit of the FFAC system tested was around 42 W/cm^2 (168 W), whilst the HLA system has an expected capacity to cool over 112.5 W/cm^2 (450 W), illustrating that liquid-based cooling can support the escalating CPU power levels, whilst traditional air-cooling cannot.
- (3) The power requirement of the two fluid movers (pump and fans) in the HLA system can greatly exceed the FFAC system, if not correctly managed.
- (4) With two fluid movers, maintaining a die temperature can be achieved by controlling either or both the pump and fan speeds, and the power cost penalty and overall COP of the system are sensitive to this choice.
- (5) By controlling the pump flow rate and fan speed, a constant setpoint die temperature can be maintained over a range of CPU heat fluxes. Importantly, this can be achieved in such a way as to minimize the cooling power and thus maximize the system COP.
- (6) For a given CPU heat flux, lowering the fan speed and increasing the liquid flow rate to achieve a die temperature setpoint minimizes cooling power and maximizes COP.

(7) For the most aggressive case investigated, with a CPU heat flux of 72 W/cm^2 (288 W) and CPU die temperature of $T_{\text{die}} = 55^\circ\text{C}$, the optimum power for the range of fluid flow rates tested is only 5.4 W, with an impressive $\text{COP} = 56$, illustrating that industry concerns about the additional energy cost of liquid-cooling can be largely addressed by engineering suitable control systems to manage the fluid movers.

The demonstrated superiority of the HLA system over the conventional FFAC system for the tested CPU heat transfer rates underscores the importance of advanced liquid-cooling solutions for energy efficient and reliable operation of modern CPUs and other processors, such as GPUs. Future research will investigate means to further increase the cooling capacity of liquid-based cooling of high-powered processors by addressing the liquid coolant temperature in an HLA system. Being limited by the ambient air temperature, including a miniature vapour-compression refrigeration system has the potential to provide sub-ambient liquid coolant temperatures, thus increasing the cooling capacity of liquid-based thermal management systems.

2.5. References

- [1] J.W. Elliott, M.T. Lebon, A.J. Robinson, Optimising integrated heat spreaders with distributed heat transfer coefficients: A case study for CPU cooling, *Case Studies in Thermal Engineering*. 38 (2022) 102354.
- [2] Intel product specifications, <https://www.intel.com/content/www/us/en/ark.html>, (accessed May 16, 2023).
- [3] M. Saini, R.L. Webb, Validation of models for air cooled plane fin heat sinks used in computer cooling, in: *ITherm 2002. Eighth Intersociety Conference on Thermal and Thermomechanical Phenomena in Electronic Systems* (Cat. No.02CH37258), 2002: pp. 243–250.
- [4] M. Saini, R.L. Webb, Heat rejection limits of air cooled plane fin heat sinks for computer cooling, *IEEE Transactions on Components and Packaging Technologies*. 26 (2003) 71–79.
- [5] A.L. Moore, L. Shi, Emerging challenges and materials for thermal management of electronics, *Materials Today*. 17 (2014) 163–174.

-
- [6] D.B. Tuckerman, R.F.W. Pease, High-performance heat sinking for VLSI, *IEEE Electron Device Letters*. 2 (1981) 126–129.
- [7] A. Mohammed Adham, N. Mohd-Ghazali, R. Ahmad, Thermal and hydrodynamic analysis of microchannel heat sinks: A review, *Renewable and Sustainable Energy Reviews*. 21 (2013) 614–622.
- [8] A.A. Alfaryjat, H.A. Mohammed, N.M. Adam, M.K.A. Ariffin, M.I. Najafabadi, Influence of geometrical parameters of hexagonal, circular, and rhombus microchannel heat sinks on the thermohydraulic characteristics, *International Communications in Heat and Mass Transfer*. 52 (2014) 121–131.
- [9] Y.J. Lee, P.K. Singh, P.S. Lee, Fluid flow and heat transfer investigations on enhanced microchannel heat sink using oblique fins with parametric study, *International Journal of Heat and Mass Transfer*. 81 (2015) 325–336.
- [10] G. Xie, J. Liu, Y. Liu, B. Sunden, W. Zhang, Comparative Study of Thermal Performance of Longitudinal and Transversal-Wavy Microchannel Heat Sinks for Electronic Cooling, *Journal of Electronic Packaging*. 135 (2013).
- [11] Y.L. Zhai, G.D. Xia, X.F. Liu, Y.F. Li, Heat transfer in the microchannels with fan-shaped reentrant cavities and different ribs based on field synergy principle and entropy generation analysis, *International Journal of Heat and Mass Transfer*. 68 (2014) 224–233.
- [12] E.G. Colgan, B. Furman, A. Gaynes, W. Graham, N. LaBianca, J.H. Magerlein, R.J. Polastre, M.B. Rothwell, R.J. Bezama, R. Choudhary, K. Marston, H. Toy, J. Wakil, J. Zitz, A practical implementation of silicon microchannel coolers for high power chips, in: *Semiconductor Thermal Measurement and Management IEEE Twenty First Annual IEEE Symposium*, 2005., 2005: pp. 1–7.
- [13] J. Zhou, X. Cao, N. Zhang, Y. Yuan, X. Zhao, D. Hardy, Micro-Channel Heat Sink: A Review, *J. Therm. Sci.* 29 (2020) 1431–1462.

-
- [14] W. Qu, I. Mudawar, Experimental and numerical study of pressure drop and heat transfer in a single-phase micro-channel heat sink, *International Journal of Heat and Mass Transfer*. 45 (2002) 2549–2565.
- [15] A.M. Kiper, Impinging water jet cooling of VLSI circuits, *International Communications in Heat and Mass Transfer*. 11 (1984) 517–526.
- [16] G.J. Michna, E.A. Browne, Y. Peles, M.K. Jensen, The effect of area ratio on microjet array heat transfer, *International Journal of Heat and Mass Transfer*. 54 (2011) 1782–1790.
- [17] A.J. Robinson, E. Schnitzler, An experimental investigation of free and submerged miniature liquid jet array impingement heat transfer, *Experimental Thermal and Fluid Science*. 32 (2007) 1–13.
- [18] A. Robinson, A Thermal–Hydraulic Comparison of Liquid Microchannel and Impinging Liquid Jet Array Heat Sinks for High-Power Electronics Cooling, *Components and Packaging Technologies, IEEE Transactions On*. 32 (2009) 347–357.
- [19] D. Nikolic, M. Hutchison, P.T. Sapin, A.J. Robinson, Hot spot targeting with a liquid impinging jet array waterblock, in: *2009 15th International Workshop on Thermal Investigations of ICs and Systems*, 2009: pp. 168–173.
- [20] A. Bhunia, C.L. Chen, On the Scalability of Liquid Microjet Array Impingement Cooling for Large Area Systems, *Journal of Heat Transfer*. 133(6): 064501 (2011)
- [21] A.J. Robinson, R. Kempers, J. Colenbrander, N. Bushnell, R. Chen, A single-phase hybrid micro heat sink using impinging micro-jet arrays and microchannels, *Applied Thermal Engineering*. 136 (2018) 408–418.
- [22] J.S. Bintoro, A. Akbarzadeh, M. Mochizuki, A closed-loop electronics cooling by implementing single-phase impinging jet and mini channels heat exchanger, *Applied Thermal Engineering*. 25 (2005) 2740–2753.

- [23] S. Liu, J. Yang, Z. Gan, X. Luo, Structural optimization of a microjet based cooling system for high power LEDs, *International Journal of Thermal Sciences*. 47 (2008) 1086–1095.
- [24] J.-Y. Chang, H.S. Park, J.I. Jo, S. Julia, A System Design of Liquid Cooling Computer Based on the Micro Cooling Technology, in: *Thermal and Thermomechanical Proceedings 10th Intersociety Conference on Phenomena in Electronics Systems*, 2006. ITherm 2006., 2006: pp. 157–160.
- [25] B.P. Whelan, R. Kempers, A.J. Robinson, A liquid-based system for CPU cooling implementing a jet array impingement waterblock and a tube array remote heat exchanger, *Applied Thermal Engineering*. 39 (2012) 86–94.
- [26] Freecor FTC | Artec Coolants, (n.d.). (accessed March 31, 2023).
- [27] R.L. Webb, Heat Exchanger Design Methodology for Electronic Heat Sinks, *Journal of Heat Transfer*. 129 (2006) 899–901.
- [28] J.R. Taylor, *An introduction to error analysis*, University Science Books, Sausalito, Calif., 1997.
- [29] R. Kempers, P. Kolodner, A. Lyons, A.J. Robinson, A high-precision apparatus for the characterization of thermal interface materials, *Review of Scientific Instruments*. 80 (2009) 095111.
- [30] R. Kempers, A. Robinson, Heated Meter Bar Techniques: What You Should Know and Why, in: *The Art of Measuring in the Thermal Sciences*, CRC Press, 2020.
- [31] Thermal Management for Intel® Xeon® Processors, Intel, <https://www.intel.com/content/www/us/en/support/articles/000006710/processors/intel-xeon-processors.html> (accessed January 20, 2023).
- [32] Sandro Villinger, How to Check and Monitor Your CPU Temperature, How to Check and Monitor Your CPU Temperature, Avast, July 2023.

- [33] Information about Temperature for Intel® Processors, Intel, <https://www.intel.com/content/www/us/en/support/articles/000005597/processors.html> (accessed May 30, 2023).

Chapter 3

An Experimental Investigation of Vapour-Compression Refrigeration Cooling and Energy Performance for CPU Thermal Management

Abstract

This case study experimentally investigates a R-134a charged Vapour-Compression Refrigeration (VCR) assisted liquid-cooling of high-powered CPUs. In this hybrid system, water-based liquid-cooling of the CPU is achieved in a pumped loop with a commercial cold plate affixed to the simulated CPU. The heat is then transferred to the evaporator heat exchanger of a miniature VCR system that rejects the energy to ambient air. Tests are performed at room temperature over a range of operating conditions, including CPU heat loads (12 W/cm^2 - 75 W/cm^2), liquid volumetric flow rates (0.5 L/min - 8 L/min), fan air speeds (2.1 m/s - 4.2 m/s), and compressor speeds (2540 rpm - 6500 rpm). Both cooling performance and additional energy consumption implications are considered. Overall, it is shown that the miniature hybrid VCR system can cool the primary liquid to below ambient temperature, even at the highest tested heat load. As a result, cooling levels that far exceed both a conventional Fan-Fin Air Cooled heat exchanger and a Hybrid Liquid-Air Cooled system are achieved. Finally, optimal combinations of liquid flow rate and compressor speed are identified for minimum energy consumption of the cooling system to maintain a setpoint CPU die temperature. It is shown that management and control of pump, compressor and fan speeds can result in significant improvement in the overall system Coefficient of Performance (COP).

3.1. Introduction

Computer processors continue to experience rapid growth in their performance, which has been driven by the increase in the number of transistors that can be fit onto an Integrated Circuit semiconductor chip. Moore's Law, which states that the number of transistors on a chip doubles every two years, has been a factor behind this trend [1]. The increase in the number of transistors has led to continuous escalation of the processing capabilities of CPUs. However, this growth in performance comes with a commensurate increase in power consumption and heat generation.

The heat generated by high-powered processors, such as the 56-core Intel Xeon Platinum 8480+ which boasts a Thermal Design Power (TDP) of 350 W, require effective thermal management systems to maintain their operating temperature within the safe limits, which is typically between 80-90°C for modern processors.

The Thermal Design Power (TDP) is defined as the maximum net power that can be dissipated by the processor cores while still maintaining a suitable operating temperature with a direct-attach Fan-Fin Air-Cooled heat sink in an ambient of 35°C [2]. However, it is important to note that the TDP is not necessarily the maximum power that can be achieved by the CPU, as it represents the power that is permitted for a particular family of air-cooled heat sinks. Indeed, higher CPU power levels can be achieved if more aggressive cooling methods are used. This is crucial in technical use scenarios such as gaming, machine learning, artificial intelligence, video rendering and other high-performance computing applications. To keep pace with the increasing power and heat fluxes of modern CPUs, the development of new and more effective cooling technologies is required.

Liquid-based cooling technology is an obvious solution for high-powered processors. Naduvilakath-Mohammed *et al.* [3], detailed an investigation of a Hybrid Liquid-Air (HLA) system that implemented a closed pumped loop of water-based liquid between a cold plate fixed to the CPU source, and a remote air-cooled heat exchanger. In terms of overall cooling performance, it was shown that this cooling method was far superior to a traditional Fan-Fin Air-Cooled heat exchanger,

facilitating lower CPU temperatures, and expanded CPU power range. Importantly, it was shown that by simultaneously controlling the pump and fan speed, the excess power consumption of the cooling system could be minimized, achieving overall system Coefficient of Performance levels commensurate with the lower performance Fan-Fin technology.

While closed loop HLA type systems, and others like it, can improve cooling performance over traditional air-cooling, they are still limited by the ambient air temperature. To improve cooling beyond this, active methods that can lower the liquid coolant to below ambient temperatures, such as thermoelectric coolers (TECs) and Vapour-Compression Refrigeration (VCR), have been considered [4]. VCR technology is the focus of this investigation, as it has been previously shown that, from both cooling performance and energy consumption perspectives, VCR technology is far superior to TEC technology [4].

The early VCR investigations of Nnanna [5] and Trutassanawin *et al.* [6] used directly attached evaporators to a simulated electronics device, and proved the concept of VCR-cooling whilst providing preliminary performance data. Nnanna [5] noted that some caution should be exercised using the direct-attach method since a portion of the evaporator cold plate contains superheated refrigerant vapour, resulting in uneven cooling of the chip. Wu *et al.* [7] designed and tested an in-house miniature VCR system for cooling PCs with a direct-attach evaporator. The work focussed on refrigerant charge and capillary tube length, and were able to tune the design to achieve 200 W of cooling of a simulated CPU. Jeng *et al.* [8] considered a hybrid-cooled direct-attach cold plate with two-phase VCR cooling focussed in the central region and single-phase $\text{Al}_2\text{O}_3/\text{water}$ nanofluid cooling the peripheral area of the heat spreader. They showed that the VCR-only cooling capacity of 270 W could be increased to 540 W with the hybrid cold plate. Poachaiyapoom *et al.* [9] conducted a study on a miniature VCR system for CPU cooling in PCs utilizing R-134a refrigerant and a direct-attach commercial microchannel CPU cold plate. Although the heat source was unrealistically large for a CPU, they achieved an impressive COP of ~ 9.0 assuming isentropic compression at its lowest compressor speed of 3000 rpm, while dissipating 200 W and maintaining the heater surface at 73.3°C . The authors noted that the above stated COPs could be as much as 27%

lower for the actual cycle. Most recently, Zhou and Li [10] investigated a miniature VCR cooling system for electronics cooling, though for a relatively low heat flux $50 \times 50 \text{ mm}^2$ heat source. For a direct-attach evaporator, the overall system weight was only 2 kg, making it potentially suitable for low heat flux portable electronics or automotive and aerospace applications. Performance was characterized over a range of ambient air temperatures up to 50°C , and showed promise for high temperature applications.

Previous research has demonstrated that VCR technology is a feasible cooling method for electronics, including CPUs, in terms of performance and compactness. VCR systems have the added benefit of being capable of supplying coolant to the heat source at temperatures below ambient, which is the main advantage over non-refrigerated liquid-cooling. However, VCR research conducted to date has focused on scenarios where the evaporator is fixed to the heat source. While this approach can provide sub-ambient temperature coolant, the convective heat transfer coefficient is limited due to presence of superheated vapour and the poor thermophysical properties of refrigerants, even with forced two-phase convection. For example, the thermal resistivity of the microchannel evaporator studied by Trutassanawin *et al.* [6] was $\sim 0.8 \text{ cm}^2\text{K/W}$, which is about double that of the liquid-cooled cold plate recently reported in Naduvilakath-Mohammed *et al.* [3]. Cooling multiple sources with this method may also pose practical mechanical and packaging challenges, such as deciding between flexible or hard piping, leak protection, and maintenance, among other issues.

Considering the above, this work investigates a hybrid VCR cooling system that implements a primary closed loop single-phase water-based liquid-cooled system, with a cold plate heat exchanger fixed to the emulated CPU. Heat is rejected from this loop to the evaporator of a secondary VCR system, which then rejects heat to the ambient air. The hybrid system thus takes advantage of the sub-ambient coolant temperature offered by VCR, together with the high convective heat transfer coefficient and overall simplicity offered by the pumped single-phase water-based coolant system.

3.2. Experimental Apparatus and Data Reduction

3.2.1. Thermal Test Vehicle with Liquid-Cooled Cold plate and Vapour-Compression Refrigeration




Figure 3.1. Schematic of the Thermal Test Vehicle (TTV) with mounted liquid-cooled cold plate.

The Thermal Test Vehicle (TTV) depicted in Figure 3.1 was fabricated to mimic the characteristics of a Central Processing Unit (CPU). The TTV was machined from an Oxygen-Free Electronic (OFE) copper block with a lower neck region of square cross-section ($A_{\text{die}} = 20 \times 20 \text{ mm}^2$) which is approximately the area of higher-powered commercial CPU dies ($\text{TDP} > 150 \text{ W}$), which are typically larger than their lower-powered counterparts (e.g. Intel i9 7980XE: $\text{TDP} 165 \text{ W}$ [11]). To simulate the heat loading of a CPU, three cartridge heaters were inserted into the 60 mm long lower neck section, capable of providing a maximum total power of 700 W. The upper section of the TTV consisted of a 2.5 mm thick overhanging section of $40 \times 40 \text{ mm}^2$ area, which represented a CPU Integrated Heat Spreader (IHS), which was heated by the lower neck section. The TTV was then placed in a 3D

printed case with insulation lining to minimize heat losses and provide a mounting fixture for the liquid-cooled cold plate.




Figure 3.2. Illustration of heated meter bar and the resulting linear Temperature vs. Distance plot.

The heat flux from the lower neck to the IHS and the temperature of the neck tip were measured using three 1.0 mm diameter T-type thermocouples. These thermocouples were evenly spaced along the length of the neck section and inserted into drilled holes to a depth of 10 mm. Figure 3.2 illustrates this setup, which enables the determination of the temperature gradient at the base of the IHS, dT/dx . This temperature gradient is used to calculate the associated heat transfer to the IHS section from,

$$q_{CPU} = -k_{Cu}A_{die} \frac{dT}{dx} \quad (3.1)$$

The temperature of the interface between the neck and IHS, T_4 , is calculated using linear extrapolation,

$$T_4 = T_3 - \left(\frac{q_{CPU}}{A_{die}} \right) \left(\frac{\Delta x_3}{k_{Cu}} \right) \quad (3.2)$$

In a CPU package, a layer of thermal interface material would exist between the die and the IHS. To account for this additional thermal resistance, a thermal resistivity value for TIM1 of $R''_{\text{TIM1}} = 0.1 \text{ K} \cdot \text{cm}^2/\text{W}$ has been assumed [12], such that $R_{\text{TIM1}} = R''_{\text{TIM1}} A_{\text{die}} = 0.025 \text{ K/W}$, and a more realistic CPU die temperature can be approximated as,

$$T_{\text{die}} \cong T_4 + q_{\text{CPU}} R_{\text{TIM1}} \quad (3.3)$$

A commercial cold plate, specifically the Nexalus Enflux™ microjet array impingement heat exchanger, was used to cool the IHS of the emulated CPU. The microjet array heat exchanger, as illustrated in Figure 3.1, is comprised of a plastic manifold that directs the flow of liquid through an array of 0.7 mm orifices, generating high velocity microjets that impact a nickel-plated copper baseplate. To ensure uniform pressure, four bolts were utilized to secure the cold plate to the TTV and were tightened using a torque wrench. A layer of Nexalus Hydronex™ Thermal Interface Material (TIM) was also applied between the base of the cold plate and the IHS in accordance with industry standards.

As depicted in Figure 3.3, the experimental setup for the VCR system has three additional areas of interest: the primary coolant (liquid) loop, the secondary coolant (VCR) loop, and the tertiary coolant (ambient air). The primary liquid coolant is a mixture of 25% Arteco-Freecor FTC® and 75% distilled water, as it provides adequate freeze protection and includes corrosion and microbial growth inhibitors [13]. To control the flow rate, a SAER CF gear pump was used in the primary loop, along with a AC10 series Parker Hannifin single-phase inverter drive to regulate the motor's torque and speed. The flow rate was measured using a high-accuracy, non-invasive UF25B Cynergy ultrasonic flow meter with an uncertainty of $\pm 3\%$ reading. To determine the pressure drop, a 26PCDFA6D Honeywell differential pressure transducer with an uncertainty of $\pm 0.25\% \text{FS}$ was used across both the cold plate and the liquid-side heat exchanger of the evaporator. Similarly, temperature drop measurements were taken across both the cold plate and the heat rejector using T-type thermocouples, which were inserted into the flow using compression fittings. The temperature probes were calibrated against a platinum RTD reference probe, resulting in an uncertainty of $\pm 0.2^\circ\text{C}$. The total penalty associated with the power

required for the liquid-side of the cooling system requires determination of the hydraulic pumping power across the cold plate and the evaporator,

$$W_{\text{pump}} = \dot{V}_{\text{liq}}(\Delta P_{\text{liq,clp}} + \Delta P_{\text{liq,e}}) \quad (3.4)$$

where $\dot{V}_{\text{liq}}$ is the volumetric flow rate of the primary liquid coolant, and $\Delta P_{\text{liq,clp}}$ and $\Delta P_{\text{liq,e}}$ is the associated pressure drop across the cold plate and the evaporator. For the HLA system depicted in Figure 3.4 and detailed in Naduvilakath-Mohammed *et al.* [3], the total hydraulic penalty is calculated identically.




Figure 3.3. Wireframe diagram of the miniature VCR-assisted liquid-cooling system.

In this case study, the heat rejection to the air was facilitated via a miniature VCR system. Refrigeration technology has yet to be widely deployed in consumer computing systems, despite their obvious potential for high-power application. It is particularly noteworthy that this VCR unit (Aspen LCM FP00039) occupies a net volume of 199 mm wide x 161 mm tall x 134 mm deep, which can be considered suitably compact for some desktops and workstations, among others computing applications. The refrigerant is R-134a (1,1,1,2-tetrafluoroethane), a hydrofluorocarbon (HFC) refrigerant commonly used in compact refrigeration systems, and constitutes the secondary coolant loop depicted in Figure 3.3. The main components of this VCR system are the compressor, condenser, throttling valve and evaporator. The latter is a 185 mm x 80 mm x 20 mm compact-braided counter flow heat exchanger to which the liquid coolant from the CPU cooling loop transfers heat, plus any energy gain from the pump to the evaporating R134a refrigerant. A 24V DC variable speed ASPEN Q9 compressor then compresses the refrigerant vapour to the air-cooled condenser, which is 195 mm x 130 mm x 26 mm and cooled by a single 120 mm FD1212-AP281EARX CeraDyna fan.

The experimental investigations presented in this thesis were conducted under steady-state operating conditions to establish the thermal-hydraulic and energetic performance limits of the investigated cooling architectures. In practical computing applications, processor heat loads are inherently transient due to variations in computational workload and the operation of processor power-management mechanisms, including dynamic frequency scaling, turbo boosting, and power-state transitions [14,15]. Modern processors can respond to workload changes on millisecond timescales, while workload execution and synchronization events may produce power fluctuations over timescales ranging from hundreds of milliseconds to several seconds[15,16]. The dynamic response of vapour-compression refrigeration systems is governed primarily by the thermal inertia of the evaporator, condenser, refrigerant inventory, and associated heat exchangers, together with the response characteristics of the compressor speed controller [17,18]. Consequently, refrigeration systems respond to changes in thermal load over finite timescales rather than instantaneously [17,18]. Furthermore, the thermal inertia of the heat exchangers generally dominates the transient behaviour of the refrigeration cycle [19]. However, the thermal capacitance of the cold plate, coolant

inventory, integrated heat spreader, and processor package provides a degree of thermal buffering that can mitigate short-duration heat-load excursions. A detailed transient investigation of the refrigeration system response to realistic processor workloads was beyond the scope of the present study and represents an important area for future research.

For conventional air-cooled and liquid-cooled systems, cooling performance is commonly quantified using the overall thermal resistance between the heat source and ambient environment. However, this metric becomes difficult to interpret for refrigeration-based cooling systems because external work is supplied to actively reduce the coolant temperature below ambient conditions. Under such circumstances, the temperature difference between the CPU die and ambient air may approach zero or become negative, resulting in thermal resistance values that are no longer representative of cooling effectiveness. To address this limitation, Huang *et al.* [20] introduced the concept of the overall thermal indicator, R'_T , though for a thermoelectric-augmented Hybrid Liquid-Air cooled electronics cooling system, which separates the contributions of the passive liquid-cooling circuit and the active refrigeration system. This approach is adopted in the present work to provide a more physically meaningful representation of system thermal performance under sub-ambient operating conditions. When adapted to the VCR system considered here, the overall thermal indicator can be defined as,

$$R'_T = \frac{T_{\text{die}} - T_{\text{ambient}}}{q_{\text{CPU}}} = R'_{\text{LCC}} + R'_{\text{VCR}} \quad (3.5)$$

which can be positive or negative. The thermal indicators for the Liquid Cooling Circuit (LCC) and Vapour-Compression Refrigeration (VCR) system are respectively,

$$R'_{\text{LCC}} = \frac{T_{\text{die}} - T_e}{q_{\text{CPU}}} \quad (3.6)$$

$$R'_{\text{VCR}} = \frac{T_e - T_{\text{ambient}}}{q_{\text{CPU}}} \quad (3.7)$$

Therefore, R'_{LCC} is a positive value quantifying the overall performance of the LCC, which includes TIM1, IHS, TIM2, cold plate, and liquid-side of the evaporator heat exchanger. Performance is improved by engineering interventions that reduce R'_{LCC} . Since $T_e < T_{ambient}$, R'_{VCR} is a negative value that reflecting the active refrigeration effect provided by the vapour-compression cycle, including the evaporator, compressor, and condenser. Increasing the magnitude of the R'_{VCR} corresponds to greater sub-ambient cooling capability.

For the VCR system, the compressor speed, and thereby the cooling capacity of the system, was varied and controlled using a motor drive board, which utilises a user supplied 0 - 5 VDC analogue signal to set the compressor speed. The firmware is designed to limit the minimum operating speed of the system to 2540 rpm, which corresponds to a voltage of 0.7 VDC or lower. Additionally, the firmware sets the maximum operating speed to 6500 rpm, which corresponds to a voltage of 4.5 VDC or higher [21]. Between the voltage range of 0.7 V and 4.5 V, there is a linear relationship between the control voltage and the compressor speed (rpm). The total power drawn by the compressor for was calculated by,

$$P_{comp} = IV \quad (3.8)$$

where I and V are the voltage and current supplied, respectively.

The VCR system is instrumented with thermocouples at the entrance and exits of both the evaporator and condenser to estimate the thermodynamic states and associated enthalpies of the refrigerant. The mass flow rate of the refrigerant, $\dot{m}_{ref}$, is not measured directly, though can be determined with reasonable accuracy by performing an energy balance at the evaporator, such that,

$$\dot{m}_{ref} = \frac{q_e}{(h_{in} - h_{out})} \quad (3.9)$$

where h is the refrigerant enthalpy. The condenser fan speed was controlled by a desktop DC power supply, and electrical fan power was determined as,

$$P_{fan} = IV \quad (3.10)$$

where I and V are the voltage and current supplied, respectively. In this study, the fan power was adjusted within a voltage range of 6 V - 15 V, while measuring the resulting air velocity using an Omega® vane anemometer with an uncertainty of $\pm 2\%$ of the reading. To determine the system's Coefficient of Performance (COP) for a given evaporator heat load, q_e , the total hydraulic power, the compressor power, as well as the total fan power, combine such that [11, 12],

$$\text{COP}_{\text{VCR}} = \frac{q_e}{W_{\text{total}}} = \frac{q_{\text{CPU}} + W_{\text{pump}}}{W_{\text{pump}} + P_{\text{comp}} + P_{\text{fan}}} \quad (3.11)$$




Figure 3.4. Wireframe diagram of the instrumented HLA cooling system.

Unlike the conventional refrigeration COP, which evaluates the thermodynamic performance of the vapour-compression cycle alone, the COP employed in this study represents a system-level performance metric that accounts

for the total electrical power consumed by the cooling infrastructure, including the compressor, pump, and fan. This definition was retained to ensure consistency with the performance assessment methodology adopted in Chapter 2 and to enable direct comparison between the FFAC, HLA, and VCR cooling architectures on a common energy-consumption basis. For the HLA system depicted in Figure 3.4 and detailed Naduvilakath-Mohammed et. al. [3], the COP is calculated by only considering the total hydraulic power and fan power in the denominator of Eq. 3.11.

3.2.2. Uncertainty Analysis

Table 3.1 lists the experimental error of the primary measurement equipment. To calculate the propagated uncertainty on derived quantities, standard procedures defined by Taylor [24] were followed, assuming that the uncertainties were independent and random. When computing the uncertainty, δZ , of a derived quantity, Z , which is dependent on measured quantities such as X , Y , and others, the uncertainty is determined based on these measured quantities as;

$$\delta Z = |Z| \sqrt{\left(\frac{\delta X}{|X|}\right)^2 + \left(\frac{\delta Y}{|Y|}\right)^2 + \dots} \quad (3.12)$$

Table 3.1. Uncertainties of measured quantities.

Measurement	Instrument	Abs. Uncertainty
Temperature, T	T-type Thermocouple	$\pm 0.2^\circ\text{C}$
Liquid Flow Rate, $\dot{V}_{\text{liq}}$	Ultrasonic Flow Meter	$\pm 3\%$ reading
Pressure Drop, ΔP_{liq}	Diff. Pressure Sensor	$\pm 0.25\%$ FS
Air Velocity, v	Vane Anemometer	$\pm 2\%$ reading
Voltage, V	Digital Multimeter	$\pm (0.09\% + 2\text{dgt})$
Current, I	Digital Multimeter	$\pm (1.0\% + 3\text{dgt})$

The rate of heat transfer through the heated meter bar system is subject to uncertainty which requires special attention. This uncertainty is not only influenced

by multiple temperature and spatial errors, but also by the regression fitting used to determine dT/dx . Additionally, this error propagates to the estimation of T_{die} , which must also be measured accurately. To address this issue, the Monte-Carlo calculation method, initially described in Kempers *et al.* [25] and later detailed in Kempers and Robinson [26], was implemented. The uncertainties associated with the rate of heat transfer for the VCR system is shown in Table 3.2

Table 3.2. Heat Transfer Uncertainties of VCR System.

Heat Transfers	Uncertainty (%)	
	Min.	Max.
T_{die} (°C)	± 0.2	± 1.9
q_{CPU} (W)	± 1.8	± 3.7
P_{comp} (W)	± 1.0	± 1.1
W_{pump} (W)	± 3.1	± 14.1
P_{fan} (W)	± 3.0	± 3.3
COP (-)	± 4.8	± 15.0

3.3. Results and Discussion

3.3.1. Effect of Heat Load and Compressor Speed

Figure 3.5 (a) & (b) shows the thermal indicator trends for different heat fluxes for fixed nominal liquid (4 L/min) and air (3.7 m/s) flows for low (2540 rpm) and high (6500 rpm) compressor speeds. For the LCC, the liquid-side convective thermal resistances in the cold plate and evaporator heat exchangers are nearly constant despite small variations in the thermophysical properties. These fixed convective thermal resistances result in the thermal indicator, R'_{LCC} , plateauing at about 0.18 K/W, with comparatively small sensitivity to heat loading and compressor speeds compared to R'_{VCR} . The small increase R'_{LCC} at lower heat loads is primarily attributed to the decreasing evaporator temperature, which is evident in Figure 3.6.

It therefore corresponds to a VCR effect rather than a thermal-hydraulic one on the LCC side.




Figure 3.5. Thermal indicator for increasing CPU heat loads at $Q_{liq} = 4$ L/min and $v = 3.7$ m/s for compressor speeds of (a) $S = 6500$ rpm, and (b) $S = 2540$ rpm.

With regards to the VCR system, Figure 3.5 also shows that for low CPU power, R'_{VCR} can become sufficiently negative that an overall negative R'_T value results. This indicates that $T_{die} < T_{ambient}$, as shown in Figure 3.6. For both compressor speeds, increasing the thermal load causes an initial steep rise in R'_{VCR} , which then tends to asymptotically approach zero. This overall trend with heat load is unsurprising since $R'_{VCR} \propto (-1/q)$ and $T_{ambient}$ is constant and higher than T_e for both compressor speeds (see Figure 3.6). However, the magnitude of R'_{VCR} , and thus the overall thermal indicator R'_T , is influenced by the VCR performance since, under a given set of operating conditions, it determines the difference $T_e - T_{ambient}$.




Figure 3.6. Cooling performance of VCR-assisted liquid-cooled system for varying heat load at $Q_{liq} = 4$ L/min and $v = 3.7$ m/s for compressor speeds of $S = 6500$ rpm, and $S = 2540$ rpm.




Figure 3.7. P-h diagram of VCR-assisted liquid-cooled system for varying heat load at $Q_{liq} = 4 \text{ L/min}$ and $v = 3.7 \text{ m/s}$ for compressor speed of (a) $S = 6500 \text{ rpm}$, and (b) $S = 2540 \text{ rpm}$.

The P-h diagrams in Figure 3.7 show the VCR thermodynamic cycle for both compressor speeds of 6500 rpm (Figure 3.7(a)) and 2540 rpm (Figure 3.7(b)). On each plot, the cycle for a low (24 W/cm^2), medium (48 W/cm^2) and high (75 W/cm^2) thermal loads are shown. As Figure 3.7 shows, increasing the heat load decreases the pressure ratio across the compressor and subsequently the evaporator pressure. This is related to the increasing heat load resulting in a proportionate increase in the refrigerant mass flow rate (see Eq. 3.9). With the fixed compressor speed and condenser pressure, the increase in mass flow rate must be accommodated by a proportionately lower specific volume of vapor at the suction side of the compressor, which requires an increase in the evaporator pressure.

From a thermal performance perspective, the end result that the increase in T_e increases with increasing heat load, as shown in Figure 3.6 for both compressor speeds. However, Figure 3.6 also shows that T_e is notably lower for the 6500 rpm case, and this results in a better overall thermal performance, as indicated by R'_{VCR} and thus R'_T being lower than that of the 2540 rpm case (see Figure 3.5). Again, this is related to the refrigerant mass flow rate being predominantly determined by the heat load. In this case, with the nearly fixed condenser pressure, lowering the compressor speed for similar refrigerant mass flow rate necessitates a lower vapor specific volume, and thus higher pressure at the suction side of the compressor. Overall, this reduces the cooling effectiveness of the overall system.

Figure 3.6 outlines the overall CPU cooling performance characteristics for varying heat load and compressor speed. It is observed that for the 6500 rpm setting, the point at which T_{die} becomes warmer than the ambient air temperature is approximately 35 W/cm^2 , whereas at the lowest rpm, it occurs at approximately 20 W/cm^2 . Also, for the 6500 rpm case, R'_{VCR} is low enough to keep the primary coolant temperatures at sub-ambient temperature, even at the highest heat load, increasing to only $\sim 8^\circ\text{C}$ at 75 W/cm^2 (300). This compressor setting can thus maintain T_{die} about 50 K below the approximate CPU throttling limit, offering sufficient potential for increasing the CPU power. Extrapolating indicates that the safe operating power limit could be as high as 470 W and throttling not occurring up to 540 W.

For the lowest compressor speed, R'_{VCR} approaches zero for higher heat loads such that $R'_T \rightarrow R'_{LCC}$ (see Figure 3.5(b)) and the cooling performance is

governed by the VCR system's capacity to maintain a sufficiently low T_e . As discussed, for a given heat load the lower compressor speed results in proportionately higher T_e , resulting in higher CPU die temperatures. In fact, despite the refrigerant temperature being sub-ambient, the primary liquid coolant temperature is marginally above ambient ($\sim 23^\circ\text{C}$) at the highest heat load. As will be discussed, this in fact does not offer any notable benefit over the HLA cooling system, though requires significant additional power to operate the cooling system.




Figure 3.8. Power penalty plots along with associated COP for low and high CPU power and compressor speeds; $Q_{\text{liq}} = 4 \text{ L/min}$ and $v = 3.7 \text{ m/s}$.

Figure 3.8 shows the power penalty and overall system COPs for low and high heat flux and the highest and lowest compressor speeds tested. It is first noted that the hydraulic power in the LCC is minimal. The pressure drop for the primary liquid coolant across the cold plate and evaporator were about the same ($\sim 16 \text{ kPa}$ and 14 kPa respectively), with only slight differences over the range of operating conditions due to temperature-viscosity variations. This results in only about 2 W of

required hydraulic power. For each test case, the compressor is the primary contributor to the power consumption of the cooling system, surpassing the combined power of the fan and pump by approximately 12X for 6500 rpm and 4X at 2540 rpm.

Considering Figure 3.6 and Figure 3.8 together, it is clear that operating the compressor at high speeds leads to improved cooling performance, with low T_{die} and reasonable system $COP = 2.2$ for the highest heat load. For critical cooling applications (e.g. high CPU power, high ambient temperatures), VCR cooling can be considered a very effective solution. However, for the 6500 rpm case and lowest heat load, this drops to $COP = 1.15$, since the overall power consumption only reduces marginally with heat load. Practically, this is excessive from an energy usage standpoint. Reducing the compressor speed to its minimum 2540 rpm increases both the $COP = 2.2$ and T_{die} , though can still be considered excessive cooling. Although the COP increase to ~ 6.4 for the low rpm case and highest heat load, the cooling performance is commensurate with what can be achieved by a well-designed HLA cooling system, which will consume considerably less power [3]. This will be discussed briefly in Section 3.3.3.

3.3.2. Effect of Primary (Liquid) and Tertiary (Air) Flow Rates

Figure 3.9 (a) illustrates the variation of the primary liquid flow rate, revealing an increase in R'_{LCC} at lower flow rates, and amounts to approximately 65% difference between the highest and lowest flow rates. However, R'_{VCR} remains unaffected by the liquid flow rate. Consequently, the evaporator temperature of the refrigerant remains approximately constant, as seen in Figure 3.10. Consequently, the observed decrease in the die temperature (T_{die}) shown in Figure 3.10 can be attributed to the corresponding increase in the liquid-side convective heat transfer coefficients within the cold plate and evaporator, rather than a VCR effect.




Figure 3.9. Thermal indicator for varying (a) primary (liquid) flow rate and (b) tertiary (air) speed for maximum tested heat flux of $q'' = 75 \text{ W/cm}^2$ and compressor speed $S = 6500 \text{ rpm}$.

Referring to Figure 3.11, the decrease in die temperature is accompanied by only a slight improvement in the total system power consumption when compared to the 4 L/min setting. Furthermore, it is noteworthy that beyond the 4 L/min flow rate, there is minimal improvement observed in R_{VCR} , R_T and T_{die} , which aligns with the findings from the HLA system investigated by Naduvilakath-Mohammed et. al. [3]. However, as Figure 3.11 shows, the pumping power escalates, having a more notable influence on the COP, which drops from 2 to 1.8.




Figure 3.10. Cooling performance of VCR-assisted liquid-cooled system for varying primary (liquid) flow rate, and varying tertiary (air) speed for maximum tested heat flux of $q'' = 75 \text{ W/cm}^2$ and compressor speed $S = 6500 \text{ rpm}$ and 2540 rpm .

Figure 3.9 (b) demonstrates that adjusting the fan speed to increase the air velocity at the condenser side has no discernible impact on any of the thermal indicators. As a result, there is no significant alteration observed in the die temperature, as shown in Figure 3.10. Despite the lack of influence on cooling

performance, decreasing the fan speed negatively affects the overall energy consumption (and the associated COP), as indicated in Figure 3.11. It was observed that as the fan speed increased from 2 to 4.2 m/s, there was a noteworthy reduction in the electrical power input to the compressor, decreasing from 176 W to 135 W. This reduction in compressor power input led to an increase COP from 1.7 to 2. The reason behind this improvement is that the low fan speed decreases the effectiveness of the condenser heat exchanger.




Figure 3.11. Power penalty plots along with associated COP for varying primary (liquid) flow rate, and varying tertiary (air) speed for maximum tested heat load (75 W/cm^2) and compressor speed (6500 rpm)

Referring to the thermodynamic cycles plotted in the P-h diagrams in Figure 3.12, it is clear that decreasing the fan speed raises the condensing temperature and associated pressure. This occurs while the evaporator temperature and pressure remain relatively constant. As a result, there is an increase in the pressure

ratio, necessitating the compressor to perform additional work on the refrigerant, consequently increasing compressor power. Increasing the fan speed further from 3.7 m/s to 4.2 m/s exhibits a similar effect, albeit less significant, likely due to diminishing increases in the air-side convective heat transfer coefficient, which limits further improvements in condenser heat exchanger effectiveness.




Figure 3.12. P-h diagram of VCR liquid-cooled system for varying condenser air velocity with fixed compressor speed, liquid flow rate and heat load.

3.3.3. Power optimization of VCR system

To illustrate the potential for optimizing the cooling energy requirement of the VCR system, a case study is performed here where the compressor and pump speeds are controlled in such a way to maintain a constant die temperature of $T_{die} = 55^\circ\text{C}$ for varying CPU heat load, similar to the analysis described in Naduvilakath-Mohammed et. al. [3]. It has already been established that operating the condenser fan speed beyond its nominal setting (3.7 m/s) only increases the overall system

fan power consumption with no improvement in cooling performance. Likewise, decreasing it below the nominal setting causes the compressor power to increase, with again no cooling performance gain. Therefore, this nominal fan speed has been kept constant in the following optimization analysis, and the heat flux is varied between 46 W/cm^2 to 83 W/cm^2 , which are the minimum and maximum heat loads that can maintain the selected setpoint die temperature with this particular cooling system setup. Primary liquid flow rates between 0.5 L/min - 8 L/min and compressor speeds between 2540 rpm - 6500 rpm are considered. It is noted here that this analysis is not exhaustive, and is provided here as an example case study to illustrate the potential energy savings associated with engineering systems that can control more than one mechanical component in the overall cooling system.

Figure 3.13 (a) & (b) plot the total cooling power consumption and associated COP to maintain die temperatures of 55°C over the range of die heat fluxes as a function of the liquid flow rate and compressor speed. Figure 3.13 (a) shows that the minimum cooling power consumption adapts to a higher liquid flow rate and lower compressor speed. For example, considering the 66 W/cm^2 case, the same overall cooling performance can be achieved with a liquid flow rate of 1.7 L/min and compressor speed of 6500 rpm , as it can with a higher liquid flow rate of 4.25 L/min and lower compressor speed of 2550 rpm . Here the total cooling power drops considerably from 155 W to 55.5 W with an associated increase in the COP from 1.7 to 4.7 , constituting a 3-fold improvement in energy performance, which is significant. The magnitude of the improvement diminishes at higher CPU heat loads since the minimum required compressor speed tends to increase. For example, considering the 78 W/cm^2 case, it is optimal to operate the VCR system with the liquid and refrigerant fluid movers set at 6 L/min and 4650 rpm respectively, resulting in an energy consumption decrease from 151 W to 109 W and COP increase from 2.13 to 3 , which is still about 40% improvement in energy performance. Overall, the die temperature can be regulated, at minimum cooling power expense, by reducing the compressor speed and increasing the liquid flow rate, and the power consumption saving can be substantial.




Figure 3.13 (a) VCR Power Consumption (PC), and (b) COP surface plots for maintaining a die temperature of $T_{\text{die}} = 55^\circ\text{C}$ for varying heat loads by varying liquid flow rate and compressor speed for $T_{\text{ambient}} = 20^\circ\text{C}$.

3.3.4. Comparison of VCR and HLA Cooling Systems

As a final note, the hybrid VCR cooling system is now compared against the HLA system described in Section 2.1. Furthermore, comparison is also made with a conventional Fan-Fin Air-Cooler (FFAC), with details of this heat exchanger and the HLA system provided in Naduvilakath-Mohammed *et al.* [3].




Figure 3.14. Cooling performance of FFAC, HLA and VCR cooling systems for varying heat load.

Figure 3.14 summarizes the cooling performance characteristic of each the FFAC, HLA and VCR systems, showing achieved T_{die} as a function of CPU heat load. For conciseness, the FFAC system is characterized for the maximum and minimum air speeds. For the HLA and VCR systems, the primary liquid coolant flow rate of 4 L/min is selected, as it has been shown both here and in Naduvilakath-Mohammed *et al.* [3] that beyond this, improvement in cooling performance is minimal. For the HLA cooling system, the maximum and minimum achievable air

speeds are included, whereas, for reasons discussed above, the air speed for the VCR system was fixed at the optimal 3.7 m/s. Results for both the maximum and minimum compressor speeds for the VCR system are considered.

The maximum power dictated by the throttling temperature for the three cooling systems are 44.2 W/cm^2 (177 W) for the FFAC system, 112.5 W/cm^2 (450 W) for the HLA system and 135 W/cm^2 (540 W) for the VCR system, with the latter two determined by extrapolation. It is noted that for the FFAC system, this is close to the stated TDP = 165 W of the Intel i9 7980XE which is being emulated here. The HLA and VCR systems are estimated to achieve 2.7 and 3.3 fold increase in potential design power respectively.

It is clear in Figure 3.14 that the VCR system at the highest compressor speed achieves the lowest die temperature of the three cooling systems tested, and does so over the full range of heat loads considered. Compared with the best performing settings for the FFAC and HLA, T_{die} is respectively 50K and 15K colder using the VCR system. However, with the VCR at its lowest compressor speed, the HLA system is the better performing of the two, especially considering much lower additional cooling power. In terms of the cooling power penalty, the FFAC and HLA require approximately 2.4 W and 7.8 W respectively. Considering the maximum heat load, this translates into $\text{COP}_{\text{FFAC}} \sim 88$ and $\text{COP}_{\text{HLA}} \sim 35$. In contrast, the power consumption of the VCR system is 143 W with associated $\text{COP}_{\text{VCR}} \sim 2.2$. Although optimization can potentially improve the COPs of both the HLA and VCR systems for a given operating temperature, this will not affect the conclusion that the main drawback of the VCR technology is the power consumption of the compressor. This being considered, future research in the field should focus on means to optimize the cooling versus energy cost relationship of miniature VCR technology, especially considering the high potential it shows for cooling very high-powered CPUs.

3.4. Conclusions

This case study aimed to investigate the cooling performance of a closed loop liquid-cooled system using Vapor Compression Refrigeration (VCR) technology for advanced thermal management of high-powered CPUs. For an emulated $20 \times 20 \text{ mm}^2$ CPU die and range of input power and heat fluxes up to 300

W and 75 W/cm^2 respectively, the experimental investigation measured key cooling performance metrics of a VCR system, such as die temperature and thermal indicator, over ranges of liquid flow rates up to 8 L/min , fan air speeds up to 4.2 m/s and compressor speeds up to 6500 rpm . The cooling system energy requirement was also assessed, including total cooling power and overall system COP. The following conclusions can be drawn;

- (1) Increasing compressor speed improves cooling performance by achieving higher compressor pressure ratios and associated lower refrigerant evaporator temperature.
- (2) The energy consumption cost is dominated by the compressor power. Higher compressor speeds may result in improved cooling performance, though at lower system COP.
- (3) Varying the primary (liquid) and tertiary (air) flow has much smaller influence on the overall cooling performance compared with compressor speed
- (4) An optimal fan speed has been determined beyond which no cooling improvement is observed for the additional fan power expense. Interestingly, reducing below the optimal has the desired effect of reducing the fan power, but causes the compressor pressure ratio to increase, resulting in higher compressor power for null gain in cooling performance.
- (5) The energy cost of the VCR cooling system can be reduced significantly by simultaneously controlling the primary liquid coolant flow rate and compressor speed. The best result was obtained at a heat load of 66 W/cm^2 , where the COP increased from 1.7 to 4.7 by managing the compressor and pump speeds, whilst maintaining the same die temperature.
- (6) Overall, sub-ambient primary liquid coolant temperatures can be achieved, even for the highest heat loads tested. As a result, the VCR system can achieve colder die temperatures compared with FFAC and HLA systems. At the highest heat load of 75 W/cm^2 (300 W), the die temperature is about 50 K lower than a FFAC heat exchanger, and 15 K lower than the HLA system tested. Although outside of the

parameter range tested here, the improved cooling performance will likely allow for adequate cooling for higher CPU powers and/or elevated ambient temperatures.

Future research will focus on developing a predictive mathematical model of the hybrid VCR system. This will include appropriate sub-models for all the system sub-components such that multi-objective optimization tools can be used to maximize cooling performance for maximum system COP.

3.5. References

- [1] M. Roser, H. Ritchie, Technological Progress, Our World in Data (2013). <https://ourworldindata.org/technological-progress> (accessed June 4, 2021).
- [2] 3rd Gen Intel® Xeon® Scalable Processors Brief, Intel (n.d.). <https://www.intel.com/content/www/us/en/products/docs/processors/xeon/3rd-gen-xeon-scalable-processors-brief.html> (accessed January 20, 2023).
- [3] F.M. Naduvilakath-Mohammed, R. Jenkins, G. Byrne, A.J. Robinson, Closed loop liquid cooling of high-powered CPUs: A case study on cooling performance and energy optimization, *Case Studies in Thermal Engineering* 50 (2023) 103472.
- [4] K. Liang, Z. Li, M. Chen, H. Jiang, Comparisons between heat pipe, thermoelectric system, and vapour compression refrigeration system for electronics cooling, *Applied Thermal Engineering* 146 (2019) 260–267.
- [5] A.G.A. Nnanna, Application of refrigeration system in electronics cooling, *Applied Thermal Engineering* 26 (2006) 18–27.
- [6] S. Trutassanawin, E.A. Groll, S.V. Garimella, L. Cremaschi, Experimental Investigation of a Miniature-Scale Refrigeration System for Electronics Cooling, *IEEE Transactions on Components and Packaging Technologies* 29 (2006) 678–687.
- [7] Z. Wu, R. Du, Design and experimental study of a miniature vapor compression refrigeration system for electronics cooling, *Applied Thermal Engineering* 31 (2011) 385–390.

- [8] L.-Y. Jeng, T.-P. Teng, Performance evaluation of a hybrid cooling system for electronic chips, *Experimental Thermal and Fluid Science* 45 (2013) 155–162.
- [9] A. Poachaiyapoom, R. Leardkun, S. Wongwises, Miniature Vapor Compression Refrigeration System for Electronics Cooling, *Case Studies in Thermal Engineering* 13 (2018) 100365. <https://doi.org/10.1016/j.csite.2018.100365>.
- [10] H. Zhou, J. Li, Development and analysis of a simple structured and economic miniature vapor compression refrigerator for cooling electronics in harsh environment, *Applied Thermal Engineering* 223 (2023) 120047.
- [11] Intel® Core™ i9-7980XE Extreme Edition Processor (24.75M Cache, up to 4.20 GHz) - Product Specifications, Intel (n.d.). <https://www.intel.com/content/www/us/en/products/sku/126699/intel-core-i97980xe-extreme-edition-processor-24-75m-cache-up-to-4-20-ghz/specifications.html> (accessed October 2, 2023).
- [12] M. Saini, R.L. Webb, Heat rejection limits of air cooled plane fin heat sinks for computer cooling, *IEEE Transactions on Components and Packaging Technologies* 26 (2003) 71–79.
- [13] Freecor FTC | Arteco Coolants, (n.d.). <https://www.arteco-coolants.com/en/products/engine-coolants/x-oat-coolants/freecor-ftc> (accessed March 31, 2023).
- [14] J.L. Hennessy, D.A. Patterson, *Computer Architecture: A Quantitative Approach*, 6th ed. Cambridge, Morgan Kaufmann, 2017.
- [15] R. Schöne, T. Ilsche, M. Bielert, A. Gocht, D. Hackenberg, Energy Efficiency Features of the Intel Skylake-SP Processor and Their Impact on Performance, in: *2019 International Conference on High Performance Computing & Simulation (HPCS)*, 2019: pp. 399–406.
- [16] D. Jensen, O. Nnorom, G. Wilkins, H. Budd, R. Rajagopal, J. Rivas-Davila, P. Levis, *EasyRider: Mitigating Power Transients in Datacenter-Scale Training Workloads*, (2026).

- [17] Y. Yao, M. Huang, J. Chen, State-space model for dynamic behavior of vapor compression liquid chiller, *International Journal of Refrigeration* 36 (2013) 2128–2147.
- [18] B. Li, A.G. Alleyne, A dynamic model of a vapor compression cycle with shut-down and start-up operations, *International Journal of Refrigeration* 33 (2010) 538–552.
- [19] M.C. Keir, A.G. Alleyne, Dynamic Modeling, Control, and Fault Detection in Vapor Compression Systems, PhD Dissertation, Air Conditioning and Refrigeration Center TR-247, College of Engineering. University of Illinois at Urbana-Champaign (2006). <https://hdl.handle.net/2142/12907> (accessed June 22, 2026).
- [20] H.-S. Huang, Y.-C. Weng, Y.-W. Chang, S.-L. Chen, M.-T. Ke, Thermoelectric water-cooling device applied to electronic equipment, *International Communications in Heat and Mass Transfer* 37 (2010) 140–146.
- [21] Liquid Chiller Module - LCM-600, Aspen Systems (n.d.). <https://aspensystems.com/products/liquid-chiller-module-2/> (accessed August 13, 2021).
- [22] R. Kempers, J. Colenbrander, W. Tan, R. Chen, A. Robinson, Experimental Characterization of a Hybrid Impinging Microjet-Microchannel Heat Sink Fabricated using High-Volume Metal Additive Manufacturing, *International Journal of Thermofluids* 5–6 (2020) 100029.
- [23] A. Robinson, W. Tan, R. Kempers, J. Colenbrander, N. Bushnell, R. Chen, A new hybrid heat sink with impinging micro-jet arrays and microchannels fabricated using high volume additive manufacturing, 2017.
- [24] J.R. Taylor, *An introduction to error analysis*, University Science Books, Sausalito, Calif., 1997.
- [25] R. Kempers, P. Kolodner, A. Lyons, A.J. Robinson, A high-precision apparatus for the characterization of thermal interface materials, *Review of Scientific Instruments* 80 (2009) 095111.

[26] R. Kempers, A. Robinson, “Heated Meter Bars: What You Should Know and Why,” Chapter 13 in *The Art of Measuring in Thermal Sciences*, Taylor and Francis (2020).

Chapter 4

Modelling and Optimization of Vapour-Compression Refrigerated Liquid-Cooling of CPUs

Abstract

The present study examines the application of miniature Vapour-Compression Refrigeration (VCR) technology for cooling of high-power Central Processing Units (CPUs). Specifically, VCR-assisted liquid-cooling, which enables sub-ambient liquid coolant temperature, is investigated in terms of cooling and energy performance. A steady-state mathematical model was developed that combines heat transfer, fluid mechanic, and thermodynamic sub-models of the compressor, condenser, expansion valve, evaporator, and liquid cooling circuit to predict CPU die temperatures under varying heat loads. The model was validated against experimental data across a wide range of operating conditions: CPU heat fluxes of 25 - 75 W/cm², compressor speeds of 2540 - 6502 rpm, air velocities of 3.3 - 4.2 m/s, and liquid coolant flow rates of 1 - 6 L/min. Subsequently, the model was integrated within a Design of Experiments (DOE) framework to investigate parameter sensitivities and perform system-level design optimization, primarily focussing on achieving target CPU die temperatures with maximum Coefficient of Performance (COP). The results and analysis indicate that the compressor speed exerts the dominant influence on cooling performance, though at the highest energy consumption cost. However, the COP can be increased significantly via design optimization which balances the cooling and power consumption trade-offs. In particular, increasing either or both the liquid coolant flow rate and evaporator size have a disproportionately high influence on the cooling performance compared with the additional energy consumption penalty, allowing CPU cooling targets to be achieved with maximum COP.

4.1. Introduction

The pursuit of higher computing capabilities has driven Central Processing Unit (CPU) manufacturers to engineer increasingly more powerful processors. Effective thermal management of the current and next generation high-performance CPUs is now a critical challenge, since adequate cooling is necessary to prevent thermal throttling or potential CPU damage. Although traditional air-cooling is low cost, it is essential to acknowledge its inherent limitations in the face of advancing CPU thermal demands [1].

Liquid-based cooling technology has emerged as a logical solution for cooling high powered processors owing to their far superior thermophysical properties compared to air. This was recently demonstrated by the present authors [2], where a closed-loop Hybrid Liquid-Air (HLA) system was investigated and directly compared with conventional fan-fin air-cooling. The HLA system featured a closed pumped loop circulating water-based coolant between a cold plate heat exchanger fixed to a CPU-type source and a remote air-cooled heat exchanger. This case study not only demonstrated close to 3-fold reduction in the source-to-sink thermal resistance of the HLA system compared with conventional fan-fin air-cooling, but it also showed that the power consumption of the HLA cooling system could be minimized by independently controlling the pump and fan speeds, to the point of being lower than the fan-fin system.

While closed-loop HLA-type cooling systems exhibit superior performance compared to conventional air-cooling, the heat source temperature is still dependent on the ambient air (sink) temperature. It is, however, possible to lower CPU temperatures even further, or equivalently achieve higher Thermal Design Power (TDP), provided technologies capable of reducing the liquid coolant temperature below ambient levels are implemented. In this study, miniature Vapour-Compression Refrigeration (VCR) is investigated as a potential technology to improve the cooling capacity of liquid-based electronics thermal management systems.

The application of Vapour-Compression Refrigeration (VCR) systems to electronic device cooling has garnered increasing research attention as power

densities continue to escalate in modern processors and high-performance computing systems (e.g. servers, workstations). Early explorations established the fundamental feasibility of direct-contact evaporative cooling for microelectronic devices. Investigations by Nnanna [3] and Trutassanawin *et al.* [4] demonstrated that compact VCR technology could sustain stable operation when directly coupled to simulated electronic heat sources. Nnanna's findings, however, underscored a critical challenge - the formation of superheated vapour pockets within portions of the evaporator, which resulted in temperature non-uniformity across the device surface which constrained the overall heat transfer potential.

Subsequent developments in system miniaturization and integration advanced the field beyond proof-of-concept demonstrations. Wu *et al.* [5] engineered a scaled-down VCR module specifically for personal computer cooling applications, systematically analysing the influence of refrigerant charge and capillary tube geometry on thermal performance. Their prototype achieved a cooling capacity of approximately 200 W for a CPU emulator, establishing key design relationships between expansion device tuning and refrigerant flow stability in miniature systems.

In pursuit of improved heat spreading and system flexibility, hybrid and multiphase cooling strategies have been investigated. Jeng *et al.* [6] proposed a dual-mode cold plate that coupled two-phase refrigerant boiling at the core region with single-phase nanofluid ($\text{Al}_2\text{O}_3/\text{water}$) cooling the periphery. This hybridization significantly enhanced the overall cooling capacity - from 270 W to 540 W - demonstrating the advantage of distributing the thermal load across multiple convective regimes. Similarly, Poachaiyapoom *et al.* [7] designed a compact R-134a-based VCR system incorporating a microchannel evaporator, achieving an impressive Coefficient of Performance of COP ~ 9 at compressor speeds of ~ 3000 rpm, while maintaining a 200 W heater at a steady-state temperature of 73.3°C . More recent efforts have shifted toward lightweight, self-contained VCR architectures suitable for portable and embedded electronic platforms. Zhou and Li [8] developed a 2 kg integrated cooling module targeting low heat flux electronics with a $50 \times 50 \text{ mm}^2$ heat source, demonstrating stable operation at ambient conditions up to 50°C . Such systems illustrate the adaptability of VCR technology

to mobile, automotive, and aerospace applications where compactness, weight, and reliability are all critical design considerations.

Collectively, these investigations establish VCR as a technically robust and energetically efficient solution for electronics cooling, particularly due to its unique capability to sustain sub-ambient coolant temperatures - a limitation inherent to conventional liquid cooling methods. Nevertheless, current implementations predominantly rely on rigid, direct-attach evaporators. This architecture constrains spatial heat transfer uniformity, as localized superheating tends to degrade two-phase heat transfer effectiveness. Furthermore, scalability to multi-chip or distributed heat sources introduces additional design complexities, including flexible refrigerant routing, mechanical reliability, and maintenance accessibility. These challenges continue to motivate research into more adaptable VCR topologies and dynamic evaporator configurations capable of maintaining uniform two-phase behaviour across varying load and geometric conditions.

To address these main limitations of direct-attach type evaporators, Naduvilakath-Mohammed *et al.* [9] proposed a system whereby the VCR cooled a single-phase liquid which was pumped in a loop between the VCR evaporator and a direct-attach cold plate heat exchanger at the emulated CPU heat source. It was found that this design could achieve as much as 20% increase in the maximum CPU power compared with the fully single-phase HLA system. However, this additional cooling came at a comparatively high energy cost due to the power consumption of the compressor. Nevertheless, it was also shown that selecting the correct settings for each the single-phase pump, the compressor, and the condenser fan, has a sizable impact on the overall fluid mover power consumption, indicating that system optimization could improve the system COP.

Although substantial progress has been made in applying VCR systems to the thermal management of electronic components, the majority of existing investigations remain experimental in nature. This represents a key limitation in the context of system design, as numerous geometric and operational parameters exert coupled influences on both thermal performance and energy consumption. Achieving sub-ambient coolant temperatures with minimal power penalties requires a systematic understanding of the parameter sensitivities and their

interdependencies. It is therefore evident that a mathematical framework is better suited for performance optimisation of VCR-assisted cooling systems, if they are to realise their full potential in advanced electronics thermal management applications.

Advancement toward such optimisation necessitates predictive mathematical modelling. However, this task is inherently complex due to the large number of design variables involved, and their non-linear interactions. A substantial body of research has advanced the understanding of component-level behaviour, system modelling, and heat exchanger performance in VCR systems. From the perspective of component-level research, the early work by Ayub [10] is an example of work intended to develop improved heat-transfer and pressure-drop correlations for plate heat exchangers used as refrigerant evaporators. Qasem and Zubair [11] likewise provided a comprehensive review of compact and microchannel heat exchangers for VCR systems, highlighting the competing effects of air-side enhancement and pressure-loss penalties. Parallel advances in compact heat exchangers, from Pettersen *et al.* [12] and Hesselgreaves *et al.* [13], also emphasize the influence of heat exchanger design (geometric compactness, fin structure, and flow arrangement) on heat-transfer effectiveness, including the potential for minimizing the refrigerant charge. The compressor is likely the most critical and complex component of the VCR system. Appropriate models must not only predict the influence of mechanical compression on the refrigerant, but do so while taking into account the volumetric, isentropic and motor electrical efficiencies [14]. As discussed by Ossorio *et al.* [15], compressor models can range from physics-based to fully empirical, with advantages and disadvantages to the selected approach. Finally, each component must be combined to create a system-level VCR model, which even in the simplest form, involves a complicated interaction of heat transfer, fluid mechanics and thermodynamics across multiple components, fluid phases and states, all which interact with the high and low temperature (source/sink) environments. Cabello *et al.* [16] provide a good example of a steady-state VCR system modelling which demonstrates the coupled behaviour of compressor, condenser, and evaporator. More recently, Zsembinski *et al.* [17] introduced a more advanced mathematical framework for predicting VCR performance. This work is particularly interesting as it clearly describes the intricacies associated with the solution algorithm of the problem.

The aim of the current study is to demonstrate, through a representative case study, what can be achieved when a structured modelling and optimisation methodology is applied to a miniature VCR-assisted liquid-cooling system for electronics thermal management. The specific experimental VCR setup used by Naduvilakath-Mohammed *et al.* [9] is mathematically modelled in conjunction with a closed-loop liquid-cooling circuit and a simulated CPU-type thermal load. Following validation against in-house experimental data, the model is integrated within a Design of Experiments (DOE) framework to illustrate its capability for assessing parameter sensitivities and identifying operating conditions that minimise energy consumption for defined cooling demands. Considering the somewhat limitless end-use scenarios where this type of system might be considered, it is clearly noted here that this case study is intended as a demonstration of methodology - showing how miniature VCR-assisted systems can be analysed and optimised - rather than as a universal solution to all electronics cooling challenges. To the best of the authors' knowledge, this represents the first reported example of a mathematically modelled and experimentally validated miniature VCR-assisted liquid-cooling system for electronics thermal management with capability for system design optimization.

4.2. Methodology

4.2.1. Model Overview

A key objective of this research is to develop a mathematical model for predicting the performance of a VCR system operating under steady state operating conditions. The model is developed for the specific case study of a hybrid miniature VCR - assisted liquid-cooling system for thermal management of high-powered CPUs, though is applicable to other applications, including but not limited to different electronic heat sources and multiple heat sources.

The refrigeration cycle is depicted in Figure 4.1. The system includes a compressor, an air-cooled condenser heat exchanger, a throttling device, an evaporator heat exchanger, a pump, and a cold plate heat exchanger which is directly attached to a CPU-type heat source. An end-to-end thermal-hydraulic-thermodynamic model involves appropriate modelling of all sub-components of the VCR system, including single-phase fluids (liquid in evaporator, air in condenser),

two-phase flow (evaporator and condenser), compressor and throttling valve, and these include variable refrigerant properties as part of the solution. The numbered points depicted in Figure 4.1 represent various states of the refrigerant throughout the cycle, as well as different states of the liquid and air in contact with the evaporator and condenser heat exchangers respectively. At the moment, the model assumes that the connecting pipes, which are generally short and well insulated, do not partake in the heat transfer and pressure drop. Figure 4.1 provides a generalized P-h diagram and brief description of various fluid states. In this study, the refrigerant is R134a.




Figure 4.1. (a) Schematic of the modelled system, (b) P-h diagram schematic of a Vapour-Compression cycle with subcooling and superheating, and (c) description of the operating points.

The efficacy of the model will be checked against the experimental measurements of Naduvilakath-Mohammed *et al.* [9]. As such, the system components are modelled to match those in the experiments e.g. a rotary compressor, a cross-flow condenser heat exchanger, a counterflow evaporator heat

exchanger, and an expansion valve. These and any other additional component sub-models will now be discussed in turn.

4.2.2. Refrigerant Cooling Loop

4.2.2.1. Compressor

An empirical correlation based, or 'black box', approach has been used in this study for the modelling of the compressor. As discussed in a recent paper by Ossorio *et al.* [15], compared to white-box (physics-based), or grey box (semi-empirical) methods, the black box approach, which relies on system specific experimental data, is flexible and straight-forward to implement whilst capturing the necessary behaviour of the system. The main drawback of this approach is that it does not provide deep physical into the compressor mechanics. Since this study is focussed on system performance, and not on compressor design, this modelling approach was deemed appropriate, and is indeed used in many other similar studies (e.g. Refs [15,17–19]). As will be discussed below, the current VCR compressor model requires three empirically derived expressions, one for the isentropic efficiency, η_{isen} , one for the volumetric displacement volume, $\dot{V}_{dis}$, and one for the refrigeration capacity, q_L .

The work power imparted to the refrigerant by the compressor is,

$$W_{comp} = \dot{m}_{ref} \cdot (h_2 - h_1) = \dot{m}_{ref} \cdot \frac{h_{2,isen} - h_1}{\eta_{isen}} \quad (4.1)$$

Here, $\dot{m}_{ref}$ is the refrigerant mass flow rate circulating through the VCR system, which is defined as,

$$\dot{m}_{ref} = \frac{\dot{V}_{dis}}{v_1} \quad (4.2)$$

where v_1 is the refrigerant specific volume on the suction-side of the compressor. Thus, to calculate W_{comp} , relations for both the volumetric displacement and the isentropic efficiency are needed. These are compressor-specific, and are here determined empirically, as discussed below.

For the volumetric displacement correlation, the method proposed by Ossorio *et al.* [15] has been adopted, as it is appropriate for variable speed compressors. Concisely, the method first requires the calculation of experimental values of the volumetric displacement over a range of operating conditions, and then a regression fit of an expression of general form $\dot{V}_{\text{dis}} = f(T_e, T_c, S)$ is performed. Here, S is the compressor speed, and T_e ($=T_6=T_7$) and T_c ($=T_3=T_4$) are the evaporator and condenser saturation temperatures, respectively. In this work, the resulting empirical correlation is,

$$\dot{V}_{\text{dis}} = (160 + 0.03S - 7.5T_c + 7.8T_e) \cdot 10^{-8} \quad (4.3)$$

which fit the experimental data with $R^2=0.95$. Although Ossorio *et al.* [15] suggested the use of higher-order and cross-correlated terms, it was found that, for the compressor used in this work, sufficient accuracy is obtained when retaining the first-order leading terms. It should be noted that this correlation takes into account variations in the compressor volumetric efficiency, as it is based on experimentally determined values. Since Eq. 4.3 is a function of three parameters, two further expressions are required for closure of the compressor model. The second compressor-specific empirical model is for the isentropic efficiency. First, the relevant enthalpies (h_1 , h_2 , $h_{2,\text{isen}}$) are determined from the experiments and isentropic efficiencies calculated as,

$$\eta_{\text{isen}} = \frac{h_{2,\text{isen}} - h_1}{h_2 - h_1} \quad (4.4)$$

Similar to above, the experimentally determined values are then regression fit, resulting in the empirical correlation,

$$\eta_{\text{isen}} = (208.8 - 1.5T_e - 37.3T_c - 0.009S) \cdot 10^{-3} \quad (4.5)$$

which fit the experimental data with $R^2 = 0.93$. The third compressor-specific empirical correlation is for what is typically termed the refrigeration capacity, q_L . Of course, $q_L = q_e$, which is the total heat transfer to the refrigerant in the evaporator *i.e.* the sum of the CPU power and hydraulic pumping power of the single-phase cooling loop. Knowing the experimental values of q_L again allows for the formulation of a straight-forward empirical correlation,

$$q_L = -1272 + 15.14T_e + 32.5T_c + 0.06S \quad (4.6)$$

which fit the experimental data with $R^2 = 0.95$.

The above expressions for $\dot{V}_{dis}$, η_{isen} , and q_L are sufficient to determine the rate of work energy imparted to the refrigerant in the compressor, W_{comp} . However, in terms of total power consumption of the compressor, P_{comp} , which is required to determine the system COP, a final relation is required since,

$$P_{comp} = \frac{W_{comp}}{\xi} \quad (4.7)$$

where ξ is the electrical efficiency of the compressor. Once again, ξ is compressor-specific and can be estimated experimentally by measuring the compressor voltage and current, such that $P_{comp} = IV$, and then taking the ratio $\xi = W_{comp} / P_{comp}$. The experimental compressor electrical efficiencies are then regression fit to generate the empirical correlation,

$$\xi = (-286.5 + 13.6T_e + 14.9T_c + 0.06S) \cdot 10^{-3} \quad (4.8)$$

which fit the experimental data with $R^2 = 0.87$.

4.2.2.2. Evaporator and Condenser Heat Exchangers

The evaporator and condenser heat exchangers constitute the principal heat transfer components of the VCR system. Both are compact finned-tube or plate-type heat exchangers and are modelled for steady-state operation under uniform and steady refrigerant and air/liquid flow. The heat transfer is modelled using a segmented ϵ -NTU framework, in which the refrigerant-side flow is divided into single-phase and two-phase regions according to its thermodynamic state. For conciseness, details of the expressions and empirical correlations employed to evaluate local Nusselt numbers, heat transfer coefficients, and friction factors are presented in Appendix A. Briefly, each heat exchanger is divided into distinct thermodynamic zones corresponding to the refrigerant phase. For example, the air-cooled condenser has three zones; (i) superheated vapour, (ii) two-phase mixture, and (iii) sub-cooled liquid. Within each zone, the refrigerant-side and air-side heat

transfer coefficient are evaluated using region-appropriate correlations to determine the local overall thermal conductance. The three zones, which constitute a parallel thermal network, are then combined to determine the overall condenser thermal conductance. Appendix A provides more detail of the sub-zone and full heat exchanger modelling, including how the ε -NTU framework is implemented.

4.2.2.2.1. Condenser Modelling




Figure 4.2. Schematic representation and geometric specifications of the cross-flow air-cooled condenser.

Figure 4.2 illustrates the air-cooled condenser heat exchanger design and dimensions used in Naduvilakath-Mohammed *et al.* [9], and thus in the current model. The heat rejected from the refrigerant flowing within the MPE tube to the forced air flowing over the sawtooth-type fins is the sum of actual work done by the compressor on the refrigerant (W_{comp}) and cooling capacity by the system (q_e), noting that q_e is the heat absorbed in the evaporator and includes the heat transfer from the CPU as well as the hydraulic power added to the liquid by the pump. The heat transfer from the condenser can thus be expressed as:

$$q_c = W_{\text{comp}} + q_e = \dot{m}_{\text{ref}}(h_2 - h_5) \quad (4.9)$$

where h_2 and h_5 are the specific enthalpies of the refrigerant entering and leaving the condenser respectively. As previously mentioned, this total heat load is

dissipated to the air via three segmented heat exchanger zones *i.e.* the superheated zone transports the portion $q_{sh,c}$, the condensing zone transports $q_{sat,c}$, and the sub-cooled zone transports $q_{sc,c}$. In this way, each sub-zone represents a pseudo-independent heat exchanger, separated by the interior two-phase region. As these three zones constitute a parallel thermal network, the total heat rejected is expressed as:

$$q_c = q_{sh,c} + q_{sat,c} + q_{sc,c} \quad (4.10)$$

$$= (UA)_{sh,c}(\Delta T_{LMTD,sh,c}) + (UA)_{sat,c}(\Delta T_{LMTD,sat,c}) + (UA)_{sc,c}(\Delta T_{LMTD,sc,c})$$

where $(UA)_{sh,c}$, $(UA)_{sat,c}$ and $(UA)_{sc,c}$ are the thermal conductance of superheated, saturated and sub-cooled regions respectively. These are determined as,

$$(UA)_{sh,c} = \frac{F_{sh}}{\left[\frac{1}{h_{air}A_{air}} + \frac{1}{h_{ref,sh,c}A_{int,c}} \right]}; \quad (UA)_{sat,c} = \frac{F_{sat}}{\left[\frac{1}{h_{air}A_{air}} + \frac{1}{h_{ref,sat,c}A_{int,c}} \right]}; \quad (UA)_{sc,c} = \frac{F_{sc}}{\left[\frac{1}{h_{air}A_{air}} + \frac{1}{h_{ref,sc,c}A_{int,c}} \right]} \quad (4.11)$$

In the single-phase regions, the refrigerant-side heat transfer coefficients, $h_{ref,sh,c}$ and $h_{ref,sc,c}$, are estimated using the modified Gnielinski correlation [20], while the condensing two-phase region is modelled using the correlation of Moser *et al.* [21]. The air-side heat transfer coefficient, h_{air} , is estimated using the appropriate empirical correlation provided by Kim and Bullard [22]. The fractional area of each region (F_{sh} , F_{sat} , and F_{sc}) is determined iteratively such that $F_{sh} + F_{sat} + F_{sc} = 1$, while satisfying overall energy balance constraints, with the two-phase region, positioned where the thermodynamic quality is $0 \leq x \leq 1$, being sandwiched between the two single-phase regions. For each region, the ϵ -NTU formulation is applied with appropriate heat capacity ratios to determine the heat transfer rate q_c and local effectiveness, ϵ . The analysis assumes negligible refrigerant pressure drop, constant specific heats within each phase region, and uniform air-side flow distribution. Detailed equations for evaluating zonal Nusselt numbers, log-mean temperature differences (LMTDs), and effectiveness relations are provided in Appendix A.

In the modelling strategy used, air speed/velocity, v , is selected as an independent variable. The air speed not only defines the air-side convective heat transfer coefficient via the Nu-Re relationship, but also the net volumetric flow rate, which further depends on the overall condenser size (length and/or number finned MPE tube rows). As will be discussed later when the model is deployed to demonstrate system-level optimization, varying the condenser size influences the system COP. In relation to the system COP, the electrical power consumption of the condenser fans, P_{fan} , must be determined as it must be added to the compressor and pump power to determine the overall fluid mover power consumption.

4.2.2.2. Evaporator Modelling




Figure 4.3. Schematic of the plate heat exchanger (evaporator) and key geometric specifications.

Figure 4.3 illustrates the liquid coolant - refrigerant evaporator heat exchanger design and dimensions used in Naduvilakath-Mohammed *et al.* [9], which is modelled in this study. The evaporator is a single-pass counter-flow plate heat exchanger (PHE) that transfers heat from the primary liquid cooling loop to the refrigerant. The analysis assumes steady-state operation, negligible pressure drop on the refrigerant side, negligible heat loss to the surroundings, negligible conduction resistance across the solid plates separating the fluids, uniform flow

distribution, and one-dimensional heat transfer perpendicular to the plate surfaces. The refrigerant enters the evaporator as a saturated two-phase mixture at low-pressure (Figure 4.1(b)) and exits as superheated vapour, which like the condenser heat exchanger discussed above, separates the evaporator into single and two-phase zones. The liquid coolant in the primary loop is a 25% by volume Freecor-water mixture.

The heat transfer at the evaporator of the system, q_e , includes the heat supplied to the liquid coolant by the CPU as well as the hydraulic power from the pump, and is related to the refrigerant-side via the energy balance,

$$q_e = \dot{m}_{\text{ref}}(h_1 - h_6) \quad (4.12)$$

where h_1 and h_6 are the outlet and inlet specific enthalpies of the refrigerant in the evaporator respectively. The overall heat transfer conductance is determined by combining the refrigerant-side and coolant-side resistances. Following the same methodology as discussed above for the condenser, two refrigerant regions are considered: two-phase region and superheated vapour region (see Figure 4.1 (b)). The total heat absorbed by the refrigerant from the liquid coolant in the evaporator is therefore given by:

$$q_e = (q_{\text{sh}} + q_{\text{sat}})_e = (UA)_{\text{sat},e}(\Delta T_{\text{LMTD},\text{sat},e}) + (UA)_{\text{sh},e}(\Delta T_{\text{LMTD},\text{sh},e}) \quad (4.13)$$

where $(UA)_{\text{sat},e}$ and $(UA)_{\text{sh},e}$ denote the overall thermal conductance of the saturated and superheated regions of the evaporator, respectively. These are determined as,

$$(UA)_{\text{sat},e} = \frac{F_{\text{sat}}}{\left[\frac{1}{h_{\text{liq}}A_{\text{liq}}} + \frac{1}{h_{\text{ref,sat},e}A_{\text{int},e}} \right]} \quad (UA)_{\text{sh},e} = \frac{F_{\text{sh}}}{\left[\frac{1}{h_{\text{liq}}A_{\text{liq}}} + \frac{1}{h_{\text{ref,sh},e}A_{\text{int},e}} \right]} \quad (4.14)$$

Here, the liquid-side single-phase heat transfer coefficient, h_{liq} , is estimated using the correlation of Okada *et al.* [23], which is suitable for water-based coolants in chevron-type PHEs. Downstream of the two-phase zone, the refrigerant becomes superheated. The single-phase heat transfer coefficient, $h_{\text{ref,sh},e}$, is evaluated using the same modified Gnielinski correlation as in the condenser analysis, and is

detailed in Appendix A. In the two-phase region, the evaporative heat transfer coefficient, $h_{\text{ref,sat,e}}$, is obtained using the correlation developed by Ayub [10] for convective boiling in plate heat exchangers. The fractional area associated with the two-phase region, F_{sat} , relates to the region where the thermodynamic quality is $x \leq 1$ and determined iteratively through convergence of the energy balance, with the fraction associated with the superheated zone satisfying the condition $(F_{\text{sh}} + F_{\text{sat}})_e = 1$. Once again, the ϵ -NTU method is applied, though here for a counter-flow heat exchanger, to calculate the overall heat transfer rate.

The liquid-side pressure drop in the evaporator is calculated by determining the friction factor [24] such that,

$$\Delta P_{\text{liq,e}} = f \frac{L}{D_H} \frac{\rho_{\text{liq}} V_{\text{liq}}^2}{2} \quad (4.15)$$

where $\Delta P_{\text{liq,e}}$ is the pressure drop (Pa), f the Darcy friction factor, L the flow length (m), D_H the hydraulic diameter (m), ρ the fluid density (kg/m^3) and V_{liq} the mean velocity (m/s).

4.2.2.3. Expansion Valve and Superheat

Assuming adiabatic conditions, the expansion process is isenthalpic such that the refrigerant expands from the condenser pressure to the evaporation pressure. In the VCR being modelled, the throttling device is a Thermostatic Expansion Valve (TXV) type, which maintains a reasonably constant degree of refrigerant superheating ($\Delta T_{\text{sh}} = T_1 - T_7$) at the exit of the evaporator. Therefore, the key expressions related to the expansion valve are;

$$h_5 = h_6 \quad (4.16)$$

$$T_1 = T_6 + \Delta T_{\text{sh}} \quad (4.17)$$

The degree of superheat required by the model was determined by regression fitting experimental data ($R^2 = 0.88$), and is predicted by,

$$\Delta T_{\text{sh}} = 7.75 + 0.25T_e - 0.41T_c + 0.0026S \quad (4.18)$$

4.2.3. Liquid Cooling Loop

To predict T_{die} at a given heat flux, the net thermal resistance between the emulated CPU source and the coolant is required. The Thermal Test Vehicle (TTV) used in the experiment being modelled, as detailed in Naduvilakath-Mohammed *et al.* [2,9], included a heat spreader, a layer of Thermal Interface Material (TIM) and the cold plate heat exchanger. A straight-forward empirical correlation for the net thermal resistance was developed from experimental data, with an $R^2 = 0.96$, and is expressed as,

$$R_{th,clp} = 0.007894 \dot{V}_{liq}^{-0.166} \left(\frac{\mu_{liq}}{k_{liq}} \right)^{-0.189} \quad (4.19)$$

where $\dot{V}_{liq}$ is the liquid coolant volumetric flow rate in m^3/s , and temperature dependence on fluid properties is captured with the viscosity to thermal conductivity ratio term. Similarly, an empirical correlation for the pressure drop across the cold plate heat exchanger was developed, with an $R^2 = 0.9$, as;

$$\Delta P_{liq,clp} = 2.45 \dot{V}_{liq}^2 \left(\frac{\mu_{liq}}{\mu_{ref}} \right)^{-0.513} \quad (4.20)$$

which shows the correct parabolic dependency on the coolant volumetric flow rate. Here, a viscosity ratio is included to account for the pressure drop sensitivity to coolant temperature, where μ_{ref} is simply a reference viscosity at $20^\circ C$.

4.2.4. System Energetics

In terms of thermal performance, the key purpose of the VCR-assisted liquid-cooling system is to maintain safe CPU operating temperatures for a specified CPU power. Thus, for a given source heat flux, T_{die} is the main thermal performance indicator. However, from an energetics standpoint, the aim of this study is to demonstrate that it is possible to minimize the energy consumption penalty associated with the cooling system. The Power Consumption (PC) penalty arises from the rate of energy consumption of the three fluid movers - the pump, the compressor, and the fans. Thus,

$$PC = W_{total} = P_{comp} + P_{fan} + W_{pump} \quad (4.21)$$

The total power consumption of the compressor, P_{comp} , was given earlier in Eq. 4.7. The pump circulates the liquid coolant between the cold plate and the evaporator, with an associated hydraulic pressure drop in each. For a given coolant flow rate, $\dot{V}_{\text{liq}}$, the total hydraulic power is thus,

$$W_{\text{pump}} = \dot{V}_{\text{liq}} (\Delta P_{\text{liq,e}} + \Delta P_{\text{liq,clp}}) \quad (4.22)$$

The total electrical fan power, $P_{\text{fan}} = IV$, along with the total compressor and pumping power, allows for the estimation of the main energetic performance indicator, the overall system Coefficient of Performance (COP);

$$\text{COP}_{\text{VCR}} = \frac{q_e}{W_{\text{total}}} = \frac{q_{\text{CPU}} + W_{\text{pump}}}{P_{\text{comp}} + P_{\text{fan}} + W_{\text{pump}}} \quad (4.23)$$

4.3. Calculation Algorithm

The VCR-assisted liquid-cooling system model is comprised of a system of coupled linear equations, and includes refrigerant in all of the sub-cooled, superheated and two-phase mixture states. To solve this, the Engineering Equation Solver (EES) software was used, as it is purpose built for solving coupled linear equations of this type, including variable thermophysical properties.

The solution algorithm is depicted in Figure 4.4 and begins by inputting the fluid mover settings (Q_{liq} , v , S), the CPU heat load (q_{CPU}), compressor electrical power (P_{comp}), pump power (W_{pump}), and ambient conditions (P_{atm} , T_{ambient}). As Figure 4.4 illustrates, the solution methodology is iterative and initialized by guessed values of the condensing and evaporating temperatures (T_c and T_e). Using these guessed values, the compressor sub-module first calculates the volumetric and isentropic efficiencies, as well as the cooling capacity. This is followed by determining the thermodynamic states of the refrigerant at each point in the cycle, including enthalpies, which are sufficient to calculate its mass flow rate ($\dot{m}_{\text{ref}}$). With this data, the condenser sub-module is then executed, where a Newton-Raphson method is used to converge the condenser heat transfer rate, q_c , by iteratively updating the condenser outlet temperature, T_5 . Once converged, the refrigerant state at the entrance of the expansion valve is updated, which revises the

thermodynamic state at the entrance to the evaporator by assuming the expansion process is isenthalpic ($h_5 = h_6$). The evaporator sub-module is then called and operates in the same manner as the condenser sub-module, though here updates the evaporator temperature, $T_6 = T_e$, until the evaporator heat transfer rate and cooling load converge. As depicted in the flowchart in Figure 4.4, the condenser and evaporator heat exchanger sub-modules are nested, with the main updated output parameter being the evaporator temperature, T_e . With this value, a final Newton-Raphson based iterative loop is entered which converges the total compressor electrical power by updating the condenser temperature, T_c . The calculations terminate when all residuals fall within the respective convergence criteria, and performance post-processing is then performed e.g. PC, COP etc.




Figure 4.4. Flowchart of the solution algorithm of the mathematical model.

4.4. Design of Experiments

The Design of Experiments (DOE) approach was implemented using the Minitab commercial software to optimize and systematically evaluate the influence of key operating and geometric parameters on the thermal and energy performance of the VCR-assisted liquid-cooling system. A fractional factorial design was first used for factor screening, followed by a central composite design approach to capture potential curvature and interaction effects between variables. The DOE model was developed to predict two principal responses, CPU die temperature (T_{die}) and net system Power Consumption (PC), under different combinations of the control and geometric design parameters. The COP is also evaluated as a response, though is necessarily maximized for a given heat source power when PC is minimized.

Table 4.1. Summary of design factors and response variables used in the DOE framework.

Category	Parameter	Unit	Range	Levels	Role/Remarks
Factors	Heat flux, q''	W/cm ²	25 – 160	5	Disturbance / Design Load
	Liquid Flow Rate, Q_{liq}	L/min	0.1 – 6.0	6	Control / Fluid Mover
	Air Velocity, v	m/s	3.3 – 4.2	4	Control / Fluid Mover
	Compressor Speed, S	rpm	2540 – 6502	6	Control / Fluid Mover
	Evaporator Length, L_e	cm	35 – 140	4	Control / Geometry
	Condenser Length, L_c	cm	87 – 348	4	Control / Geometry
Responses	Die Temperature, T_{die}	°C	—	—	Thermal Control Target
	Power Consumption, PC	W	—	—	System Electrical Load
	Coefficient of Performance, COP	—	—	—	Energy Efficiency Indicator

Six independent variables were considered: CPU heat flux (q''), liquid coolant flow rate (Q_{liq}), air velocity (v), compressor speed (S), evaporator length (L_e), and condenser length (L_c). The selected ranges and levels shown in Table 4.1 reflect the experimental test range in Naduvilakath-Mohammed *et al.* [9] for Q_{liq} , v and S (the fluid movers), and expanded ranges for the source heat flux, evaporator

length and condenser length. The latter two are the main geometric design parameters evaluated for the case study discussed later since, from a practical perspective, it is straight-forward to select and install different heat exchangers for different end-use applications. It should be emphasized here that this is not a limitation of the design process, as each individual component in the system, including internal geometric design parameters etc., can be varied in the developed predictive model. However, the aim here is to demonstrate that the methodology of mathematical modelling of the VCR-assisted liquid-cooling system, when combined with a structured DOE approach, can be used as a tool to not only interrogate and analyse parameter sensitivity, but also perform system design optimization. Specifically in this study, it will be demonstrated that operational control and component design parameters (factors) could be determined which result in optimized system performance (responses) - typically by minimizing energy consumption (maximizing COP) while achieving a target cooling performance. Furthermore, the parameter ranges selected represent a wide range of potential end-use cooling scenarios - from low to high thermal loads, from compact to larger system size - ensuring that the DOE framework captures the system-level trade-offs between thermal control and energy efficiency over a broad range of possible design requirements.

4.5. Results and Discussions

4.5.1. Model Validation

In order to ensure confidence in the use of the VCR-assisted liquid-cooling system mathematical model presented in Section 4.2, its accuracy is confirmed by comparing predictions against the experimental data reported in the author's previous publication [9]. The model is configured to reproduce identical settings and component dimensions of the experiments, with parameter variations covering the associated operating conditions: heat flux 25 - 75 W/cm², compressor speed 2540 - 6502 rpm, air speed 3.3 - 4.2 m/s, and liquid flow rate 1 - 6 L/min. The predicted and measured system temperatures for the VCR-assisted liquid-cooling system are compared in Figure 4.5, showing consistently close agreement over a wide range of operating conditions. The data points for all monitored locations (T_1 - T_9 and T_{die})

fall predominantly within the $\pm 5\%$ deviation bands, confirming the reliability of the mathematical model predictions.

To quantify the level of agreement, the Mean Absolute Relative Deviation (MARD) and Mean Relative Deviation (MRD) are calculated for each temperature location, as defined by,

$$\text{MARD} = \frac{1}{n} \sum |(T_{\text{exp}} - T_{\text{mod}})/T_{\text{exp}}| \times 100\%$$

$$\text{MRD} = \frac{1}{n} \sum (T_{\text{exp}} - T_{\text{mod}})/T_{\text{exp}} \times 100\%$$




Figure 4.5. Comparison of predicted and measured system temperatures over a wide range of heat source powers, compressor speeds, air speeds, and liquid coolant flow rates.

The calculated results are summarized in Table 4.2. The MARD values range from 0.3% to 1.5%, while MRD values vary between -1.02% and +0.9%, indicating minimal bias and excellent overall agreement between predicted and experimental

temperatures. Importantly, the CPU die temperature (T_{die}) shows a deviation of only 0.5%, confirming the model's strong predictive capability for the emulated CPU. Overall, the average deviation remains within $\pm 1\%$, validating the efficacy of the VCR-assisted liquid-cooling system mathematical model and justifying its integration with the DOE methodology for optimization and parameter sensitivity analysis.

Table 4.2. Quantitative validation of the VCR-assisted liquid-cooling system mathematical model showing the Mean Relative Deviation (MRD) and Mean Absolute Relative Deviation (MARD) for each temperature measurement point.

	T_1	T_2	T_5	T_6	T_8	T_9	T_{die}
MRD (%)	-0.01	-1.02	-0.9	-0.3	0.2	0.9	0.37
MARD (%)	0.3	1.5	1.3	0.5	0.5	0.9	0.5

4.5.2. Validation of the DOE-Based Optimization Framework

Following the successful validation of the VCR-assisted liquid-cooling system mathematical model against experimental data, the next required step involves confirmation of the accuracy of the DOE-based optimization framework. An experimental dataset was created in Naduvilakath-Mohammed *et al.* [9] that allowed for the interpolation/extrapolation of the system's Power Consumption (PC) and Coefficient of Performance (COP) across a range of source heat fluxes, compressor speeds, air flow rates, and liquid coolant flow rates. The analysis was then used in a test case scenario whereby a target of 55°C was selected for the die temperature, and the fluid-mover settings that yielded the minimum PC and maximum COP were calculated. The process was repeated for a range of source heat fluxes, ranging between 46 - 86 W/cm^2 . Selected results from Ref [9] are listed in Table 4.3.

To validate the DOE described in Section 4.4 against the experimentally derived values listed in Table 4.3, first the lengths of the evaporator and condenser were selected as '*hold values*' factors, such that L_e and L_c were fixed and equal to those of the experimental apparatus. Next, a heat flux factor was selected, and also defined as a *hold value*. Finally, the die temperature response goal was selected as

a target at $T_{\text{die}} = 55^{\circ}\text{C}$, whereas the PC response goal was selected to be minimized. This was repeated for each heat flux factor listed in Table 4.3. Thus, the main outputs of this DOE optimization process are, for the selected heat flux, the fluid mover settings which result in the output responses PC_{min} (and thus COP_{min}) for the selected target die temperature.

Table 4.3. Comparison of experimentally derived and mathematically predicted optimum operating parameters from the DOE-based optimization for varying heat loads and target $T_{\text{die}} = 55^{\circ}\text{C}$ with ambient air at $T_{\text{ambient}} = 20^{\circ}\text{C}$.

Heatflux (W/cm ²)	Compressor Speed (rpm)			Liquid Flowrate (L/min)			Air Velocity (m/s)			Power Consumption (W)			Coefficient of Performance (-)		
	exp.	num.	ΔS (%)	exp.	num.	ΔQ_{liq} (%)	exp.	num.	Δv (%)	exp.	num.	ΔPC (%)	exp.	num.	ΔCOP (%)
46	2550	2540	-0.4	0.1	0.1	0.0	3.7	3.6	-2.7	46.4	49.7	7.1	4.12	3.96	-3.9
60	2550	2698	5.8	3.0	3.3	9.3	3.7	3.5	-5.4	49.3	52.8	7.0	5.06	4.97	-1.7
78	4650	4600	-1.1	6.0	5.9	-1.6	3.7	3.5	-5.4	107.8	102.1	-5.3	3.01	3.14	4.3
86	6500	6178	-4.9	5.5	5.9	7.2	3.7	3.6	-2.7	140.3	130.9	-6.4	2.55	2.62	2.7

Table 4.3 presents the comparison between experimentally calculated fluid mover settings and the associated PC_{min} & COP_{min} against the those predicted by the DOE-based optimization approach for the range of source heat fluxes. The model-DOE predicted fluid mover settings are seen to closely match the experimentally determined ones, being within $\pm 10\%$. In terms of the experimental and DOE-predicted PC_{min} and COP_{min} , it follows that similarly good agreement is observed, since these are dependent on the fluid mover and heat flux settings. Importantly, the predicted fluid mover settings that minimize PC follow closely the experimental trends with increased heat flux. This will be discussed in more detail in the next section, though it can be concisely said that Naduvilakath-Mohammed *et al.* [9] not only concluded that a nominal air speed of ~ 3.7 m/s is suited across all test conditions, but also that the PC_{min} is achieved when the compressor speed is as low as feasible and the liquid coolant flow rate then set as high as necessary to achieve the required cooling effectiveness. The DOE results are consistent with these earlier observed behaviours, noting first the relatively constant air speed ranging between 3.5 - 3.7 m/s. Also, for the lowest heat flux, the DOE predicted compressor speed is at the minimum of the range. When the heat flux is increased

to 60 W/cm^2 , the compressor speed remains near the minimum and only the liquid coolant flow rate is increased to accommodate the additional required cooling. At the higher heat loads, the liquid coolant flow rate is indeed at or close to the maximum. As will be discussed, this is because, relative to the compressor, the sensitivity of T_{die} to Q_{liq} is comparatively high contrasted to the sensitivity of PC.

The developed VCR-assisted liquid-cooling system mathematical model has been implemented into a DOE and then validated against experimental data. Additionally, the combined model-DOE approach then independently reproduced experimentally determined optimization values and trends, thus providing a second level of validation by confirming that the mathematical prediction approach not only predicts the thermal-hydraulic-thermodynamic behaviour accurately, but also captures the integrated performance trends essential for design optimization. This agreement confirms that the DOE-integrated mathematical model reliably predicts this multi-parameter problem and can be confidently used for predictive design optimization *i.e.* identifying operating and component design configurations that achieve desired thermal targets with minimal energy input.

4.5.3. Parameter Sensitivity Analysis

A parameter sensitivity analysis using the DOE has been performed to systematically evaluate how changes in input variables (factors) impact the model's output (responses). Specifically, a local sensitivity study was conducted within the DOE framework to further analyse the parameter sensitivity of the operating and design factors (heat flux (q''), liquid flow rate (Q_{liq}), air speed (v), compressor speed (S), and heat exchanger lengths (L_e , L_c), on the key responses (die temperature, (T_{die}), Power Consumption (PC), and Coefficient of Performance (COP). The local sensitivity analysis results are shown in Figure 4.6, and indicate the following;

1. As expected, the source heat flux has the strongest influence (+50%) on increasing T_{die} . It also has a significant influence ($\sim +37\%$) on increasing PC, and this is related to the fact that higher cooling loads increase the required refrigerant mass flow rate, resulting in additional compressor work for a given compressor speed [9].

2. In terms of the fluid movers, the compressor speed is by far the most influential factor, exerting the highest sensitivity on PC (+41.9%) and the highest influence on reducing T_{die} (−27.6%), confirming its strong effect on both energy and thermal performance. Compared to the compressor, the liquid coolant flow rate has a smaller, yet still important influence on reducing T_{die} (−8.7%). However, compared with the compressor, it has a disproportionately small influence on PC (+9.2%), which is consistent with the observations described earlier, that liquid coolant-side control offers substantial temperature regulation with comparatively small energy penalty. These results reinforce the conclusion of Naduvilakath-Mohammed *et al.* [9] that, compared increasing compressor speed, increasing the coolant flow rate is the better option in terms of managing the trade-offs between cooling and power consumption. Furthermore, T_{die} shows small sensitivity to the condenser air speed over the range tested, though a proportionately small yet non-negligible influence on PC.




Figure 4.6. Local sensitivity of system responses (T_{die} , PC/COP) to operating and design factors (q'' , Q_{liq} , v , S , Le , Lc).

3. In terms of the evaporator and condenser heat exchanger sizes, increasing Le shows the second highest influence on reducing T_{die} (−11.1%). Similar to the liquid-coolant flow rate, this strong cooling sensitivity is associated with a

comparatively minor influence on increasing PC (+3.7%). One would then expect that L_e will be an important factor in terms of achieving T_{die} targets with minimum power consumption. Conversely, T_{die} is insensitive to L_c , though some sensitivity to decreasing PC (-4.8%) is observed. In fact, L_c is the only factor that has the effect of reducing the net fluid mover power consumption.




Figure 4.7. Response surface plots of the DOE-predicted die temperature (left) and power consumption (right) showing the combined effects of (a-b) liquid coolant and compressor speed, and (c-d) evaporator and condenser size.

Response surface plots are shown in Figure 4.7(a) and (b) for a fixed cooling load, air speed and heat exchanger size, to further emphasize the push-pull relationships that coolant flow rate and compressor speed have on the T_{die} - PC trade-off. Since the cooling load strongly influences the refrigerant mass flow rate, increasing S has the effect of lowering the compressor suction-side pressure, and thus saturation temperature in the evaporator [9]. This monotonically reduces the temperature of the liquid coolant exiting the evaporator heat exchanger, and thus lowers T_{die} (by means of lowering the liquid coolant temperature), and this is indicated in Figure 4.7 (a). However, Figure 4.7 (b) shows that improving cooling by

increasing S has associated with it a very steep rise in PC , which is due to the increased compressor work associated with the escalating pressure ratio. On the other hand, Figure 4.7 (a) highlights that, for a given compressor speed, the lowest T_{die} occurs at the highest liquid coolant flow rate, and this due to the reduced thermal resistance of the cold plate heat exchanger. Notably, the minimum T_{die} in Figure 4.7 (a) is positioned where both Q_{liq} and S are at their maximum settings. Conversely, Figure 4.7 (b) shows that for a given compressor speed, PC is not strongly influenced by Q_{liq} , reiterating that the T_{die} - PC trade-off is best managed by setting the liquid flow to maximum, and subsequently reducing the compressor speed until the target cooling level is achieved.

In a similar fashion, response surface plots are shown in Figure 4.7(c) and (d) for a fixed cooling load and fluid mover settings, and emphasize that L_e has much more pronounced influence on decreasing T_{die} than it does on increasing PC . As will be discussed later, increasing L_e increases the evaporator heat exchanger effectiveness resulting in colder liquid-coolant entering the cold plate, and thus lower CPU temperature. However, the associated frictional pressure drop is also higher, resulting in a small increase in pumping power and thus PC . Conversely, Figure 4.7(c) shows that L_c has little influence on T_{die} yet a small, but not negligible, influence on PC . Although increasing the condenser adds some additional fan power (for a given air speed), increasing its length also, like the evaporator, has the effect increasing the heat exchanger effectiveness. All other factors being equal, a higher condenser effectiveness has the obvious effect of reducing the condenser temperature and associated pressure. This reduces the required pressure lift and associated power consumption of the compressor. Overall, from a cooling and power consumption trade-off perspective, one would anticipate from Figure 4.7 that larger heat exchangers are better in terms of maximum cooling and minimum power consumption, though this of course negatively impacts the compactness of the overall cooling system.

4.5.4. Design Optimization Case Study

The above discussion has demonstrated that certain design settings and parameters, such as liquid flow rate, compressor speed, and evaporator length, have a positive influence in terms of cooling of the CPU. However, increasing any

of the fluid mover settings required additional power, which is detrimental to the system COP. Furthermore, these effects are not de-coupled, which makes it challenging to determine bespoke designs and operational settings which achieve cooling targets whilst maximizing system COP. This being said, with the DOE-integrated mathematical model developed here, the design process becomes straight forward, and indeed the design cycle time very short.

Table 4.4. DOE optimization case study test matrix.

Case	q''	S	Q_{liq}	v	L_e	L_c	T_{die}	PC	Optimization
0	HV	UC	UC	UC	HV	HV	TG	Min	SOO
1	HV	UC	UC	UC	UC	UC	TG	Min	SOO
2	HV	UC	UC	UC	Min	Min	TG	Min	MOO
3	Max	UC	UC	UC	UC	UC	TG	Min	MOO

HV=Hold Value, TG=Target Goal, UC=Unconstrained, Max=Maximize Goal, Min=Minimize Goal, SOO=Single-Objective Optimization, MOO=Multi-Objective Optimization.

To illustrate this, the DOE framework will now be applied to demonstrate its utility for system design optimization. Table 4.4 gives an overview of the selected demonstration cases, showing a mix of different design scenarios, including single and multi-objective design problems. The design objective(s) of the selected case studies are as follows;

Case 0: For a selected high heat flux, low target die temperature, unconstrained fluid mover settings, and fixed evaporator and condenser sizes, minimize PC (maximize COP). This case represents the baseline commercial system tested in Naduvilakath-Mohammed *et al.* [9].

Case 1: For a selected high heat flux and low target die temperature, as well as unconstrained fluid mover and evaporator and condenser size settings, minimize PC. This case demonstrates how drop-in replacement heat

exchangers can influence the COP compared with the commercial system of Case 0.

Case 2: For a selected high heat flux, low target die temperature, and unconstrained fluid mover settings, simultaneously minimize PC, L_e and L_c . This case demonstrates how drop-in replacement heat exchangers can influence the COP and compactness compared to the off-the-shelf commercial system of Case 0.

Case 3: For a selected high target die temperature (now chosen as the approximate throttling temperature of modern CPUs), and unconstrained fluid mover settings, simultaneously maximize the CPU heat flux whilst minimizing PC. This case demonstrates the highest cooling load potential of the refrigerated cooling system at the highest achievable COP.

Table 4.5 summarises the control factors and system responses for all four optimisation demonstration cases. Each case represents a distinct objective within the DOE framework, combining fixed and goal-driven variables to explore trade-offs among energy efficiency, compactness, and thermal performance. Of particular note, in every test case where power consumption (PC) is set as a minimisation objective, the liquid coolant flow rate remains at its maximum setting. This confirms earlier findings that higher liquid flow enhances evaporator-side heat transfer and allows cooling targets to be achieved with reduced compressor workload - an energetically efficient means of sustaining low die temperatures.

From a system energetics perspective, Table 4.5 reveals the benefits of case-by-case optimization. Relative to the baseline (Case 0), the energy-efficiency prioritized configuration (Case 1) achieves the same cooling performance with less than half the compressor power ($131 \rightarrow 55.6$ W), resulting in more than two-fold increase in COP ($2.7 \rightarrow 6.5$). The energy-efficiency and compactness prioritized design case (Case 2) achieves the cooling performance of the baseline case with a marginal ($\sim 28\%$) increase in the required heat exchanger sizes, though with reduced power consumption ($131 \rightarrow 111.2$ W) and higher COP ($2.7 \rightarrow 3.2$), reflecting the expected efficiency versus size trade-off. Under the extreme heat-load condition for the maximum allowable T_{die} (Case 3), the compressor speed and heat

exchanger sizes are maximized to achieve a near doubling of the heat flux, with only 6% increase in PC relative to the baseline case (131→ 137.9 W), thus achieving a reasonably high COP = 4.8. Overall, these design optimization case studies highlight how coordinated variation of compressor speed, coolant flow rate, and heat exchanger sized enable targeted optimisation of miniature VCR-assisted liquid-cooling systems for either maximum efficiency, compactness, or cooling capacity, thus demonstrating the flexibility and utility of the DOE-based design framework.

Table 4.5. DOE optimization case study results.

Case	Objective	Factors					Responses			
		Heat Flux, q'' (W/cm ²)	Compressor Speed, S (rpm)	Liquid Flow Rate, Q_{liq} (L/min)	Air Velocity, v (m/s)	Length (cm)		Die Temperature, T_{die} (°C)	Power Consumption, PC (W)	Coefficient of Performance, COP (-)
						Evaporator, L_e	Condenser, L_c			
0	Baseline	86 [Hold Value]	6178	5.9	3.6	70 [Hold Value]	174 [Hold Value]	55 [Goal: Target]	131 [Goal: Minimize]	2.7 [Goal: Maximize]
1	Energy Efficiency	86 [Hold Value]	3197	6.0	4.2	140	348	55 [Goal: Target]	55.6 [Goal: Minimize]	6.5 [Goal: Maximize]
2	Compact Design	86 [Hold Value]	5167	6.0	4.2	90.3 [Goal: Minimize]	222.1 [Goal: Minimize]	55 [Goal: Target]	111.2 [Goal: Minimize]	3.2 [Goal: Maximize]
3	Extreme Heat Load	156 [Goal: Maximize]	6498	6.0	4.2	140	348	100 [Goal: Target]	137.9 [Goal: Minimize]	4.7 [Goal: Maximize]

Figure 4.8 and 4.9 compare the thermodynamic cycles and selected performance indicators of the three demonstration cases (0, 1 and 2) with identical heat flux and die temperature targets ($q'' = 86$ W/cm², $T_{die} = 55^\circ\text{C}$). The corresponding control factors and responses are summarised in Table 4.5. In order to better understand how Case 1 can achieve the same CPU cooling performance as the baseline case with a 140 % increase in COP, it is first noted that each case employs the maximum liquid coolant flow rate (~6 L/min) as expected. This being the case, the thermal resistance of the cold plate heat exchanger remains essentially the same (Eq. 4.20). Thus, to achieve the same T_{die} for the same CPU heat flux, the liquid outlet temperature from the evaporator (T_9) must be nearly

identical in both Cases 0 & 1, and this is confirmed in Figure 4.9(a), where $T_9 \approx 13^\circ\text{C}$ for each. Considering the associated cycles on the P-h diagram in Figure 4.8, it is evident that the larger evaporator of Case 1 operates at a higher evaporating pressure. As Figure 4.9 (a) shows, the associated refrigerant saturation temperature (T_6) required to achieve the required $T_9 \approx 13^\circ\text{C}$ is thus higher for Case 1. In terms of the implication on the power consumption, Case 1 thus has the lower pressure ratio, and as observed in Table 4.5, demands a lower compressor speed to generate the required pressure lift. This finding is consistent with the parameter sensitivity analysis discussed earlier, whereby the compressor speed showed, by a large margin, the highest impact on PC.




Figure 4.8. P-h diagram comparing all demonstration cases (0, 1, 2) with $q'' = 86 \text{ W/cm}^2$ and $T_{\text{die}} = 55^\circ\text{C}$.




Figure 4.9. (a) Evaporator performance parameters (T_9 , T_6 , ϵ) and (b) power consumption budget breakdowns for demonstration cases (0, 1, 2) for $q'' = 86$ W/cm² and $T_{\text{die}} = 55^\circ\text{C}$.

From a heat transfer point of view, the reason why Case 1 can reach $T_g \approx 13^\circ\text{C}$ at the higher evaporator temperature (T_6) is due to evaporator heat exchanger effectiveness, ε . As shown in Figure 4.9(a), ε is close to unity for Case 1 (larger evaporator), compared with $\varepsilon \approx 0.54$ for Case 0 (smaller evaporator). Additionally, the larger condenser also has a higher effectiveness which reduces the average temperature of the condenser and associated condenser pressure (Figure 4.8). This also plays positively into the pressure ratio and compressor power reduction, comparable to that observed in the evaporator. To emphasize the overall impact, the power consumption budgets for Cases 0 & 1 are given in Figure 4.9 (b). Although both the pump and fan powers increase for Case 1 due to the larger heat exchanger sizes compared to the baseline, these are more than offset by the reduction in the compressor power, illustrating that for a given cooling requirement, proper sizing of heat exchangers combined with optimal flow control can increase the system COP significantly.

Though increasing the heat exchanger sizes is beneficial in the context of system energetics, the main trade-off is in relation to overall system compactness. To address this, the multi-objective optimisation (MOO) Case 2, that simultaneously minimises PC, L_e and L_c , is considered. Although not the smallest possible configuration to achieve $T_{die} = 55^\circ\text{C}$, it represents a balanced compact design with $L_e = 90.3$ cm and $L_c = 222.1$ cm, approximately ~29% larger than the baseline. In the same manner as Case 1, Figure 4.8 and 4.9 indicate that Case 2 maintains the same coolant outlet temperature ($T_g \approx 13^\circ\text{C}$) with higher evaporator temperature/pressure, and thus compressor speed/power, due to the increase in evaporator effectiveness $\varepsilon \approx 0.72$. The resulting COP = 3.2 reflects a 18.5% energy efficiency improvement relative to the baseline, achieved with a moderate 29% increase in total heat exchanger size. As a point of note, additional simulations were carried out where only the heat exchanger sizes were minimized, and resulted in substantially more compact design, with $L_e = 51.9$ cm, $L_c = 90.1$ cm (~30% smaller than the baseline commercial system). Unsurprisingly, all the fluid movers were set to near maximum capacity, resulting in a lower COP = 2.5.

In summary, for the same thermal load and target cooling performance, Figure 4.8 and 4.9 collectively demonstrate that enlarging the heat exchanger

surface area, particularly on the evaporator side, reduces the pressure ratio and associated compressor work, thus producing notable gains in COP. Importantly, integrating the mathematical model with a DOE-based optimization strategy demonstrates that this design framework can balance multiple potential design objectives, including system compactness and energy use minimization.

A final MOO demonstration case is now considered where the CPU power is maximized for a target die temperature of $T_{\text{die}} = 100^{\circ}\text{C}$ (approximate throttling temperature of CPUs) whilst the power consumption is minimized. This is Case 3 in Table 4.5, and is representative of what can be considered a challenging cooling scenario (high heat flux). In the following it is compared relative to Case 1 since the heat exchanger size factors are similarly unconstrained during the optimization process (Table 4.4), PC is also minimized, and it can be considered a similarly posed scenario of challenging.

Table 4.5 shows that the only parameter that is different between Cases 1 and 3 is the compressor speed, which for Case 3 is approaching, but not at, the maximum of the studied range. Considering the above discussions, this is expected since the power consumption objective is being minimized. Thus, the liquid flow rate and heat exchanger sizes are at their respective allowable maximums to balance the cooling and power consumption trade-offs. Overall, by running the compressor at close to its highest speed and targeting the higher throttling temperature, a 1.8-fold increase in CPU power ($344\text{ W} \rightarrow 624\text{ W}$) results.

Referring to the cycles on the P-h diagram in Figure 4.10(a), the main influence is that increasing the compressor speed results in an increase in the evaporator pressure, thus reducing the pressure ratio ($2.2 \rightarrow 1.6$). Even though the cooling load is almost 2-times larger, the enthalpy differences across the evaporator, $h_1 - h_6$, are nearly identical, meaning that the refrigerant mass flow rates is also approximately doubled for Case 3. The higher refrigerant flow rate is enabled by the elevated compressor speed ($3197 \rightarrow 6498\text{ rpm}$) combined with the higher vapour density at the compressor inlet.





Figure 4.10. (a) P-h diagram and (b) Power Consumption (PC) budget breakdowns comparing Cases 1 & 3.

In terms of system energetics, Figure 4.10(b) shows the expected compressor power increase of Case 3 ($15.1 \rightarrow 97.3$ W), whilst the fan and pump power remain the same. Despite this approximate tripling of the compressor power, only a 28% reduction in the COP ($6.5 \rightarrow 4.7$) results. Thus, these two cases, which are considered here as challenging in terms of cooling performance, result in very similar designs and operational strategies to optimize energetic performance, with the only key difference being the compressor speed. From a practical perspective this is interesting since, from a control's perspective, it implies that the most energetically efficient means of maintaining a die temperature over a range of CPU compute loads (*i.e.* heat loads), is to vary only the compressor speed.

4.6. Conclusions

This work investigated the thermal and energetic performance of a miniature Vapour-Compression Refrigeration (VCR) assisted liquid-cooling system for high-power CPUs. A predictive mathematical model was developed, validated against experimental measurements, and then integrated within a structured Design of Experiments (DOE) framework to investigate parameter sensitivities and perform system-level optimization. The main findings are as follows:

- (1) The full source-to-sink mathematical model of the miniature VCR-assisted liquid-cooling system predicted experimental measurements with acceptable accuracy over a wide range of operating conditions, thus confirming its efficacy as a design tool.
- (2) Integration of the validated model with a structured DOE methodology accurately reproduced experimentally determined settings and trends for optimal energetic efficiency, thereby confirming the framework's capability to predict system-level performance and identify optimal operating points.
- (3) Parameter sensitivity analysis revealed that compressor speed exerted the strongest influence on both cooling performance and power consumption. However, excessive compressor power can be mitigated by performing system-level optimization to maximize the COP. It was determined that appropriate sizing of the VCR heat exchangers combined with correct

balancing of the fluid mover settings can significantly increase the COP for a given CPU target temperature and heat load.

- (4) In terms of balancing the cooling and power consumption trade-offs, increasing either or both the liquid coolant flow rate and evaporator size were determined to have disproportionately high influence on the cooling performance compared with the additional energy consumption penalty, allowing CPU cooling targets to be achieved with maximum COP. The condenser air speed and size exhibited smaller but non-negligible effects, primarily influencing overall system power consumption.
- (5) A selection of DOE-based case studies demonstrated the flexibility and utility of the developed mathematical framework for diverse design and optimization scenarios. Coordinated variation of operating and geometric parameters enabled over 2-fold improvement in the system COP compared with the baseline off-the-shelf commercial VCR system, while maintaining cooling performance targets. This highlights how targeted optimization can balance trade-offs among compactness, cooling effectiveness, and energy efficiency.

The main potential technical shortcoming of the present work relates to the assumption of negligible refrigerant-side pressure drops across the evaporator and condenser heat exchangers. This simplification was necessary due to the experimental system's lack of in-stream pressure measurements, particularly across the compressor. As the system performance is highly sensitive to compressor characteristics, future work incorporating pressure instrumentation will likely improve model accuracy. Furthermore, this investigation focused on cooling and energy performance. Future studies will explore heat recovery opportunities from the VCR condenser using higher compression ratio compressors and liquid-cooled condensers to evaluate the broader potential of VCR-assisted liquid-cooling systems for sustainable thermal management applications.

4.7. References

- [1] M. Saini, R.L. Webb, Heat rejection limits of air cooled plane fin heat sinks for computer cooling, *IEEE Transactions on Components and Packaging Technologies* 26 (2003) 71–79.
- [2] F.M. Naduvilakath-Mohammed, R. Jenkins, G. Byrne, A.J. Robinson, Closed loop liquid cooling of high-powered CPUs: A case study on cooling performance and energy optimization, *Case Studies in Thermal Engineering* 50 (2023) 103472.
- [3] A.G.A. Nnanna, Application of refrigeration system in electronics cooling, *Applied Thermal Engineering* 26 (2006) 18–27.
- [4] S. Trutassanawin, E.A. Groll, S.V. Garimella, L. Cremaschi, Experimental Investigation of a Miniature-Scale Refrigeration System for Electronics Cooling, *IEEE Transactions on Components and Packaging Technologies* 29 (2006) 678–687.
- [5] Z. Wu, R. Du, Design and experimental study of a miniature vapor compression refrigeration system for electronics cooling, *Applied Thermal Engineering* 31 (2011) 385–390.
- [6] L.-Y. Jeng, T.-P. Teng, Performance evaluation of a hybrid cooling system for electronic chips, *Experimental Thermal and Fluid Science* 45 (2013) 155–162.
- [7] A. Poachaiyapoom, R. Leardkun, S. Wongwises, Miniature Vapor Compression Refrigeration System for Electronics Cooling, *Case Studies in Thermal Engineering* 13 (2018) 100365.
- [8] H. Zhou, J. Li, Development and analysis of a simple structured and economic miniature vapor compression refrigerator for cooling electronics in harsh environment, *Applied Thermal Engineering* 223 (2023) 120047.
- [9] F.M. Naduvilakath-Mohammed, M. Lebon, G. Byrne, A.J. Robinson, An experimental investigation of vapor compression refrigeration cooling and

- energy performance for CPU thermal management, *Case Studies in Thermal Engineering* 54 (2024) 103963.
- [10] Z.H. Ayub, Plate Heat Exchanger Literature Survey and New Heat Transfer and Pressure Drop Correlations for Refrigerant Evaporators, *Heat Transfer Engineering* (2010).
- [11] N.A.A. Qasem, S.M. Zubair, Compact and microchannel heat exchangers: A comprehensive review of air-side friction factor and heat transfer correlations, *Energy Conversion and Management* 173 (2018) 555–601.
- [12] J. Pettersen, A. Hafner, G. Skaugen, H. Rekstad, Development of compact heat exchangers for CO₂ air-conditioning systems, *International Journal of Refrigeration* 21 (1998) 180–193.
- [13] J. Hesselgreaves, The Impact of Compact Heat Exchangers on Refrigeration Technology and CFC Replacement, *International Refrigeration and Air Conditioning Conference* (1990).
- [14] J. Gonçalves, C. Hermes, C. Melo, J. Barbosa Jr, Experimental Mapping of the Thermodynamic Losses in Vapor Compression Refrigeration Systems, *Journal of the Brazilian Society of Mechanical Sciences* 24 (2011) 159–165.
- [15] R. Ossorio, J. Marchante-Avellaneda, E. Navarro-Peris, Development of compact empirical models for variable-speed compressors for the prediction of energy consumption, mass flow and discharge temperature, *Applied Thermal Engineering* 244 (2024) 122666.
- [16] R. Cabello, J. Navarro, E. Torrella, Simplified steady-state modelling of a single stage Vapour-Compression plant. Model development and validation, *Applied Thermal Engineering* 25 (2005) 1740–1752.
- [17] G. Zsembinski, A. de Gracia, P. Moreno, R. Rovira, M. Ángel González, L.F. Cabeza, A novel numerical methodology for modelling simple Vapour-Compression refrigeration system, *Applied Thermal Engineering* 115 (2017) 188–200.

- [18] H. Wan, T. Cao, Y. Hwang, S.-D. Chang, Y.-J. Yoon, Machine-learning-based compressor models: A case study for variable refrigerant flow systems, *International Journal of Refrigeration* 123 (2021) 23–33.
- [19] J. Ma, X. Ding, W.T. Horton, D. Ziviani, Development of an automated compressor performance mapping using artificial neural network and multiple compressor technologies, *International Journal of Refrigeration* 120 (2020) 66–80.
- [20] T.M. Adams, M.F. Dowling, S.I. Abdel-Khalik, S.M. Jeter, Applicability of traditional turbulent single-phase forced convection correlations to non-circular microchannels, *International Journal of Heat and Mass Transfer* 42 (1999) 4411–4415.
- [21] K.W. Moser, R.L. Webb, B. Na, A New Equivalent Reynolds Number Model for Condensation in Smooth Tubes, *Journal of Heat Transfer* 120 (1998) 410–417.
- [22] M.-H. Kim, C.W. Bullard, Air-side thermal hydraulic performance of multi-louvered fin aluminum heat exchangers, *International Journal of Refrigeration* 25 (2002) 390–400.
- [23] Okada, K., Ono, M., Tomimura, T., Okuma, T., Konno, H., & Ohtani, S. (1972). Design And Heat Transfer Characteristics of New Plate Heat Exchanger. *Heat Transfer - Japanese Research*, 1, 90-95.
- [24] R.K. Shah, D.P. Sekulic, *Fundamentals of Heat Exchanger Design*, John Wiley & Sons, 2003.
- [25] Y.A. Çengel, *Heat and Mass Transfer: A Practical Approach*, McGraw-Hill, 2007.
- [26] T.L. Bergman, F.P. Incropera, A.S. Lavine, D.P. DeWitt, *Introduction to Heat Transfer*, John Wiley & Sons, 2011.

Chapter 5

Modelling and Optimization of a Dual-Purpose Vapour-Compression System for CPU Liquid-Cooling and High Temperature Heat Recovery

Abstract

This work investigates a dual-purpose vapour-compression assisted liquid-cooling system capable of simultaneously providing adequate CPU cooling whilst upgrading waste heat temperatures to $\geq 80^{\circ}\text{C}$, which is comparable to temperatures reported for solid-sorbent based Direct Air Capture (DAC); a technology that extracts CO_2 directly from ambient air. A steady-state mathematical model is developed for a liquid (source)-to-air (sink) vapour-compression assisted liquid-cooling system which is subsequently validated against in-house experiments across a wide range of operating conditions (75 - 135 W/cm^2 CPU heat flux, 2174 - 4500 rpm compressor speeds, 0.9 - 1.1 m/s air velocity, and 1 - 6 L/min liquid coolant flow rate). The validated mathematical model is then modified to a heat-recovery configuration by implementing a refrigerant-to-water condenser heat exchanger within the model. The augmented mathematical model is then integrated within a Design of Experiments (DOE) methodology which is subsequently applied to a set of representative case studies, demonstrating a systematic design methodology that ensures safe CPU die temperatures while delivering heat recovery fluid temperatures suitable for DAC, concurrently optimising overall system Coefficient of Performance (COP). Results show that properly sized evaporator and condenser heat exchangers, combined with coordinated control of coolant flow and compressor speeds, can achieve CPU thermal targets while delivering $\geq 80^{\circ}\text{C}$ heat recovery fluid, with optimized heat pump COPs ranging between 3.8 and 5.9. The study demonstrates a potential approach for achieving DAC-relevant heat recovery temperatures, supporting future data-centre decarbonisation efforts.

5.1. Introduction

Global data centre power consumption continues to rise as digital infrastructure expands, particularly with the rapid advancement in AI and high-performance computing. Although facility-level energy efficiency improvements of conventional air-cooled data centres could, in the past, offset some of the growing energy use, recent studies indicate that computing demand is now rapidly outpacing these gains, contributing to escalation in net data centre energy use [1]. Further to this, conventional air-cooling is becoming progressively less feasible to manage the power and heat fluxes generated by contemporary CPU/GPU server and rack architectures [2,3]. Consequently, the data centre industry is shifting to liquid-cooling, which provides superior heat transfer performance, lower auxiliary power use, with the overall aim to increase data centre compute density with lower PUE [4].

Most of the electrical power drawn by the computational hardware is dissipated as heat. As such, data centres constitute a dense and continuous source of recoverable thermal energy. Several studies have indeed shown that integrating this heat into district heating networks can significantly reduce fossil fuel consumption and associated CO₂ emissions [5]. This trend is reinforced by recent EU regulation -including the revised Energy Efficiency Directive (2023/1791) and the EU Data Centre Sustainability Reporting Framework [6–8] - which focussed on systematic monitoring and utilisation of waste heat. However, using data centre waste heat in district heating is constrained, not only by geographic proximity of the data centre, but also by fluctuations in both the seasonal heating demands and diurnal mismatch, thus limiting annual net utilisation of data centre waste heat. These have seeded interest in *round-the-year* heat recovery applications, which can not only balance the supply and demand issue, but also can be deployed regardless of the geographic location of the data centre.

A high potential *round-the-year* heat recovery end-use application for data centre waste heat is Direct Air Capture (DAC) [9], which is a CO₂ removal technology. DAC facilities operate by pulling the CO₂ directly from the atmosphere and implements a heat-triggered chemical process to separate it. As recently

discussed by Díaz-Marín *et al.* [10], data centre waste heat is a high potential source for powering solid-sorbent DAC technology. However, solid-sorbent based regeneration requires heat energy at temperatures in the region of 80 - 100°C for desorption [10]. As Marín *et al.* [10] note, if data centre waste heat can be supplied at the required desorption temperatures, the integration with DAC systems have the potential to deliver far greater climate benefit than district heating alone.

Air-cooled data centres cannot reliably provide heat recovery temperatures higher than ~35 - 45°C [11–13], and are thus not suitable for DAC. Prior attempts to extract heat from air-cooled data centre computing infrastructure require air-to-water heat pumps with large temperature lifts, resulting in low COP and poor system-level energetics [14]. Air-cooling inefficiencies, including recirculation and nonuniform airflow, further limit heat-recovery potential [15]. Overall, from both the perspective of data centre cooling and as a means for heat recovery in DAC applications, air-cooling of data centres is impractical.

Direct Liquid Cooling (DLC) refers to the method by which a cold plate heat exchanger is fixed to the CPU/GPU processor. Single-phase DLC is the case where single-phase liquid, typically water-based, is used as the coolant. The thermophysical properties of water-based coolants in DLC applications are such that the cooling intensity is so high that a process termed Hot Water Cooling (HWC) can be implemented. HWC refers to the scenario where the liquid cold plate cooling intensity is sufficiently high that ‘hot water’, meaning coolant supply temperatures in the 40-65°C range, can still maintain safe processor temperatures [16]. Research has reported coolant outlet temperatures of 60 - 70°C with DLC-HWC, together with substantial reductions in mechanical cooling energy consumption resulting in heat recovery efficiencies above 80% [17–20]. These temperatures are suitable for district heating and sorption refrigeration [21], however, they remain insufficient for DAC. Furthermore, as CPU/GPU processor power and heat flux levels continue to escalate, it is most likely that this temperature deficit will increase *i.e.* required coolant inlet/outlet temperatures will decrease, moving the heat recovery temperature further from what is required for DAC.

In anticipation of processor cooling limits associated with DLC of processors, Vapour-Compression Refrigeration (VCR) has been researched as a practical

means of generating sub-ambient coolant temperatures in cold plate heat exchangers. Several studies have investigated miniature systems for VCR-integrated electronics cooling, focussing exclusively on refrigeration and associated cooling and energy performance [22–29]. In a recent study by Naduvilakath-Mohammed *et al.* [25], VCR-assisted liquid-cooling of high-powered CPUs was investigated experimentally. The study showed that, compared to conventional DLC, supplying sub-ambient (refrigerated) liquid coolant to the cold plate decreased the emulated CPU die temperature by 15°C- 25°C over a heat flux range between 12 W/cm² - 75 W/cm², thus giving headroom for increased CPU power. Naduvilakath-Mohammed *et al.* [26] furthered this investigation with a detailed mathematical model of the VCR-assisted liquid-cooling system in combination with a Design of Experiments (DOE) approach to demonstrate a design methodology whereby full system optimization can achieve given CPU thermal targets with minimized additional energy consumption, achieving COPs as high as 6.5. However, in the context of heat recovery, these prior works that focused exclusively on CPU cooling, typically maintained condenser temperatures of only about 40°C- 50°C [25–29], which are insufficient for DAC.

In the context of utilizing data centre waste heat in DAC applications, the key thermal challenge is to maintain safe processor operating temperatures whilst simultaneously producing sufficiently high heat recovery fluid temperatures, and these are at odds with one another. A further challenge is associated with managing the additional power consumption of the cooling system. Considering these, this study proposes and investigates a dual-purpose miniature vapour-compression assisted liquid coolant system capable of being integrated directly into data centre servers. The system acts as both a refrigeration system for adequate CPU cooling and a heat pump for DAC-relevant heat recovery temperatures. Consistent with the approach developed by Naduvilakath-Mohammed *et al.* [26], a validated steady-state mathematical model of the system is integrated with a DOE framework to identify how compressor operation, heat exchanger sizing and coolant flow settings influence system performance. The DOE method is then used in specific case study examples to demonstrate a design process that maintains safe CPU die temperatures while producing DAC-suitable heat recovery fluid temperature - all whilst simultaneously maximizing system COP. To the best of knowledge, no

published work has yet demonstrated optimized miniature vapour-compression assisted liquid-cooling capable of maintaining sufficiently low liquid coolant temperatures for safe CPU operation, while simultaneously producing hot water at the $\geq 80^{\circ}\text{C}$ required for DAC applications.

5.2. Experimental Apparatus

The experimental facility used in this study is similar to that used in the previous study by Naduvilakath-Mohammed *et al.* [25]. The key differences between the two experiments are associated with the compressor. In Naduvilakath-Mohammed *et al.* [25], the research was solely focussed on energetically efficient CPU cooling, and as such a highly efficient and low pressure ratio ($\sim 1\text{-}5$) compressor was most suitable. However, this generates commensurately low condenser temperatures, in the range of $\sim 40^{\circ}\text{C}$ - 50°C . This study, which further aims to investigate the potential of high-temperature heat recovery, the compressor was selected to operate at higher pressure ratios ($\sim 6\text{-}10$) in order to maintain low evaporator temperatures (circa 5°C - 20°C) required for CPU cooling, with high condenser temperatures (circa 75°C - 80°C), as established by the elevated temperature heat recovery fluid. A second key difference between this experiment and that of Naduvilakath-Mohammed *et al.* [26] is related to modelling of the compressor. The VCR system used in Naduvilakath-Mohammed *et al.* [25] did not facilitate in-stream pressure measurements across the compressor, and this had implications with respect to the mathematical modelling accuracy. Conversely, the facility used in this work was easily retrofitted with pressure sensors, thus facilitating more accurate representation of the compressor in terms of developing the empirical correlations for mathematical modelling of the system. Other than these, the two experimental facilities are otherwise functionally the same, so only a concise description is given here, and the reader is directed to Ref [25] for more in-depth detail.

As depicted in Figure 5.1, the system consists of three interconnected sub-systems: the Thermal Test Vehicle (TTV), the primary single-phase liquid coolant loop, and the secondary Vapour-Compression Refrigeration (VCR) loop. The copper TTV is constructed to mimic a typical Central Processing Unit (CPU),

including a $20 \times 20 \text{ mm}^2$ section to represent the footprint of the die, and a 2.5 mm thick overhanging section of $40 \times 40 \text{ mm}^2$ area to replicate the Integrated Heat Spreader (IHS). Power is supplied to the TTV via three cartridge heaters capable of providing a maximum total power of 700 W. Three 1.0 mm diameter T-type thermocouples are inserted into the lower neck section of the TTV to measure both the heat flux as well as the temperature at the base of the IHS [25,30]. Consistent with previous studies [25,26,30], a Nexalus Enflux™ cold plate heat exchanger was mounted rigidly to the TTV to cool the IHS of the emulated CPU, with a layer of Nexalus Hydronex™ Thermal Interface Material (TIM) to improve thermal contact.




Figure 5.1. Wireframe diagram of the miniature VCR assisted liquid-cooling system.

In the primary single-phase loop, the liquid is a 25:75 volumetric mixture of Arteco Freecor FTC® and distilled water, which circulates through both the cold plate heat exchanger fixed to the TTV, and the compact brazed-plate evaporator heat exchanger of the VCR unit. The coolant is circulated by two Nexalus NXP-20 pumps connected in series, enabling flow rates of up to ~6 L/min. The flow rate is measured using non-invasive UF25B Cynergy ultrasonic flowmeter with $\pm 3\%$ uncertainty, while differential pressure across both the cold plate and the evaporator are measured with Honeywell 26PCBFA6D transducers with $\pm 0.25\%$ uncertainty. The total hydraulic power of the primary liquid coolant loop is then simply the sum of those across the cold plate and evaporator heat exchangers, such that,

$$W_{\text{pump}} = \dot{V}_{\text{liq}}(\Delta P_{\text{liq,e}} + \Delta P_{\text{liq,clp}}) \quad (5.1)$$

where $\dot{V}_{\text{liq}}$ is the primary liquid flow rate, and $\Delta P_{\text{liq,clp}}$ and $\Delta P_{\text{liq,e}}$ are the pressure drops across the cold plate and evaporator respectively.

The secondary VCR loop is a modified Rigid DV3220E-P module, with net outer dimensions of 330 mm wide x 200 mm tall x 180 mm deep, and is charged with R-134a. The main components of this VCR unit are the compressor, expansion device (capillary tube-type) and condenser and evaporator heat exchangers. The compressor is a 24 V variable-speed Rigid QX3202VDL with a displacement of 3.25 cc. The compressor speed is controlled by an electronic motor drive board that accepts a 1-5 VDC analogue signal. The firmware constrains the operating range between approximately 2177 rpm at 1 V and 4500 rpm at 5 V, and the relationship between control voltage and compressor speed across this interval is approximately linear. The total compressor power is simply,

$$P_{\text{comp}} = IV \quad (5.2)$$

where I and V are the voltage and current supplied, respectively. Two WIKA 32 bar pressure gauges are positioned across the compressor to measure the inlet, P_e , and outlet, P_c , pressures, which will be used in a later section when empirically characterizing the compressor. It is also worth mentioning here that thermocouples are positioned at each the evaporator inlet/outlet and condenser inlet/outlet, which

combined with the associated pressure measurements, facilitate estimation of the thermodynamic states of the refrigerant throughout the cycle.

More detail about both the condenser and evaporator heat exchangers will be provided in latter mathematical modelling sections. Concisely, the condenser is 295 mm × 125 mm × 38 mm in size and is an air-cooled finned flat-tube type heat exchanger. Two axial fans provide airflow over the condenser exterior surface, and the air velocity is measured using an Omega vane anemometer with $\pm 2\%$ uncertainty. The condenser fan voltage is varied up to 13 V with a benchtop DC power supply, and like the compressor (Eq. 5.2), the electrical fan power, P_{fan} , is calculated as the product of the supply voltage and current. The evaporator is a compact brazed counter-flow heat exchanger with outer dimensions of 120 mm × 55 mm × 27 mm. The heat transfer rate to the refrigerant in the evaporator, q_e , often termed the cooling capacity, is the sum of the heat transfer from the TTV, q_{CPU} , and the hydraulic power to the coolant from the pump, W_{pump} (Eq. 5.1), such that,

$$q_e = q_{\text{CPU}} + W_{\text{pump}} \quad (5.3)$$

These together allow for the calculation of the overall system Coefficient of Performance (COP) [26],

$$\text{COP} = \frac{q_e}{W_{\text{pump}} + P_{\text{comp}} + P_{\text{fan}}} \quad (5.4)$$

Table 5.1 summarizes the uncertainties associated with the primary measurement instruments, whereas Table 5.2 gives the propagated uncertainty for all derived quantities, which follow the standard methodology outlined by Taylor [31], assuming that individual measurement errors are independent and random [25].

Table 5.1. Uncertainties of measured quantities.

Measurement	Instrument	Abs. Uncertainty
Temperature, T	T-type Thermocouple	$\pm 0.2^{\circ}\text{C}$
Liquid Flow Rate, $\dot{V}_{\text{liq}}$	Ultrasonic Flow Meter	$\pm 3\%$ reading
Pressure Drop, ΔP_{liq}	Diff. Pressure Sensor	$\pm 0.25\%$ FS
Air Velocity, v	Vane Anemometer	$\pm 2\%$ reading
Refrigerant Pressure	Analogue Pressure Gauge	$\pm 1.6\%$ FS
Voltage, V	Digital Multimeter	$\pm (0.09\% + 2\text{dgt})$
Current, I	Digital Multimeter	$\pm (1.0\% + 3\text{dgt})$

Table 5.2. Heat transfer uncertainties of VCR system.

Heat Transfers	Uncertainty (%)	
	Min.	Max.
$T_{\text{die}} (^{\circ}\text{C})$	± 0.2	± 1.9
$q_{\text{CPU}} (\text{W})$	± 1.8	± 3.7
$W_{\text{pump}} (\text{W})$	± 3.1	± 14.1
$P_{\text{fan}} (\text{W})$	± 3.0	± 3.3
$P_{\text{comp}} (\text{W})$	± 1.0	± 1.1
COP (-)	± 4.8	± 15.0

5.3. Mathematical Modelling

5.3.1. Model Overview

This study develops a steady-state mathematical model of a miniature VCR assisted liquid-cooling system for CPU thermal management with heat recovery. The general modelling framework builds upon that previously detailed in Naduvilakath-Mohammed *et al.* [26]. Indeed, the overarching mathematical modelling strategy in this study is the same. Therefore, to avoid unnecessary

overlap with Naduvilakath-Mohammed *et al.* [26], the discussion below will focus primarily on relevant aspects that differentiate the two studies. With this in mind, two things should be noted:

1. In the context of building and validating the mathematical model, the VCR system considered here uses different constituent hardware components (compressor, evaporator, condenser, throttling device), which requires different mathematical models, though they are integrated and solved in the same way as in Naduvilakath-Mohammed *et al.* [26].
2. In this study, a key objective is heat recovery with high heat recovery fluid temperatures ($> 80^{\circ}\text{C}$). As such, two mathematical models are presented. The first model, heretofore referred to as the Validation Model, is as the name suggest solely for validating against experimental data. This model configuration is depicted in Figure 5.2(a) and all system components are modelled in accordance with the experimental facility described above. Since the air-cooled condenser is not practical for heat recovery applications, a second model, heretofore referred to as the Heat Recovery Model, is developed where two of the VCR components are different. First and foremost, the condenser is modelled as a water-cooled compact braised heat exchanger, which is practical in terms of absorbing and transporting the recovered heat from the system. For simplicity, the same heat exchanger type used in the evaporator is modelled. Secondly, the capillary tube expansion device is modelled as a Thermostatic Expansion Valve (TXV) in order to avoid excessive superheating of the refrigerant at the compressor inlet. This model configuration is depicted in Figure 5.2(b).

As there are multiple fluids with changing properties to keep track of, Figure 5.2(c) is included which depicts the labelled Vapour-Compression thermodynamic cycle as well as a list of associated operating points of the three fluids, which can be referenced to Figure 5.2(a) and (b).




Figure 5.2. (a) Schematic of *Validation Model* configuration, (b) schematic of *Heat Recovery Model* configuration (note: the Heat Recovery Unit and Pump are shown for illustration purposes and are not modelled specifically here), and (c) P-h diagram schematic the Vapour-Compression thermodynamic cycle with description of the relevant operating points

5.3.2. Refrigerant Cooling Loop

5.3.2.1. Compressor

As already mentioned, the compressor used in this study is different to that of the Naduvilakath-Mohammed *et al.* [26] investigation. Specifically, the compressor selected for this study operates at considerably higher pressure ratios in order to achieve higher condenser pressures and temperature. Furthermore, in this new test facility, the suction and discharge pressures are measured directly,

enabling pressure-based formulations of the compressor empirical correlations, opposed to the temperature-based ones used previously [26]. This modelling approach is supported by the recent literature, where pressure-based models have been shown to reduce correlation complexity while improving accuracy [32,33]. It has also been reported that incorporating pressure data is particularly important for high pressure ratio compressors [34–37].

The power imparted to the refrigerant by the compressor is,

$$W_{\text{comp}} = \dot{m}_{\text{ref}} \cdot (h_2 - h_1) = \dot{m}_{\text{ref}} \cdot \frac{h_{2,\text{isen}} - h_1}{\eta_{\text{isen}}} \quad (5.5)$$

where h_i is the specific enthalpy associated with the state (isen = isentropic), $\eta_{\text{isen}} = (h_{2,\text{isen}} - h_1)/(h_2 - h_1)$ is the compressor isentropic efficiency, and $\dot{m}_{\text{ref}}$ is the refrigerant mass flow rate. Thus, to calculate W_{comp} , relations for both the refrigerant mass flow rate and isentropic efficiency are required.

For correlating the refrigerant mass flow rate, the method proposed by Marchante-Avellaneda *et al.* [32,33] and Ossorio *et al.* [38] has been adopted, where experimental mass flow rate data is regression fit to an empirical relation of the form $\dot{m}_{\text{ref}} = f(P_e, P_c, S)$. The resulting correlation is,

$$\dot{m}_{\text{ref}} = (255 - 0.0041S - 93.2P_e + 1.74P_c + 6.12P_eP_c) \cdot 10^{-5} \quad (5.6)$$

which fits the experimental data with $R^2 = 0.89$. In a similar manner, the experimentally determined isentropic efficiency data is regression fit to the correlation,

$$\eta_{\text{isen}} = 210.25 + 0.00234S - 26.3P_e - 8.8P_c + 2.03P_eP_c \quad (5.7)$$

which fits the data with an adequate $R^2 = 0.86$.

The third compressor-specific correlation required is that of the refrigeration/cooling capacity q_L , which of course is equal to the evaporator heat load q_e (Eq. 5.3). Using the experimentally determined values with $q_L = q_e$, the empirical correlation is,

$$q_L = -441.2 + 0.055S + 264.7P_e + 14.96P_c - 5.44P_eP_c \quad (5.8)$$

which fits the experimental data with $R^2 = 0.95$.

The correlations for $\dot{m}_{\text{ref}}$, η_{isen} , and q_L are sufficient to fully model the refrigerant-side thermodynamics of the compressor, including non-isentropic behaviour. However, the rate of work supplied to the refrigerant, W_{comp} , is less than the net power supplied to the compressor owing to motor inefficiencies. Defining the compressor electrical efficiency as $\xi = W_{\text{comp}} / P_{\text{comp}}$, a final empirical correlation is developed of the form,

$$\xi = 1.99 - 0.00021S - 0.68P_e + 0.0057P_c + 0.024P_eP_c \quad (5.9)$$

which fit the experimental data with $R^2 = 0.83$.

5.3.2.2. Evaporator and Condenser Heat Exchangers

The heat exchangers modelled in the present study are depicted in Figure 5.3 and Figure 5.4. It is reminded that the air-cooled condenser of 4 is only used in the *Validation Model*, whereas the plate heat exchanger in Figure 5.3 is the same for both the evaporator and condenser in the *Heat Recovery Model*.




Figure 5.3. Schematic of the plate heat exchanger and key geometric specifications. This heat exchanger design is the evaporator in both the Validation Model and Heat Recovery Model, as well as the condenser in the Heat Recovery Model.




Figure 5.4. Schematic representation and geometric specifications of the cross-flow air-cooled condenser used in the *Validation Model*.

Within the VCR system, the evaporator and condenser heat exchangers have both single and two-phase regions, which require a zone-based ε -NTU modelling strategy, as detailed in Naduvilakath-Mohammed *et al.* [26]. Succinctly, for a fluid-to-fluid heat exchanger with N zones, with each transferring an amount of heat q_i , the total rate of heat transfer in the heat exchanger, q , is the sum,

$$q = \sum_{i=1}^N q_i = \sum_{i=1}^N (UA)_i (\Delta T_{LMTD,i}) = \sum_{i=1}^N \left\{ \frac{F_i}{\left[\frac{1}{h_{i,1}A_{i,1}} + \frac{1}{h_{i,2}A_{i,2}} \right]} \right\} (\Delta T_{LMTD,i}) \quad 5.10$$

where i can represent a superheated, subcooled or two-phase zone. The $(UA)_i$ terms represent the zonal-specific overall thermal conductance. As shown, this not only depends on the local thermal resistances of each fluid, $R_{i,j}=1/(h_{i,j}A_{i,j})$, but also the fraction, F_i , each zone occupies, noting that $\sum_{i=1}^N F_i = 1$. Thus, a key requirement for evaporator and condenser modelling is the appropriate prediction of the convective heat transfer coefficients for each fluid, $h_{i,1}$ and $h_{i,2}$ in the specific zone. The heat transfer coefficient prediction methods used for both single and two-phase flow, as well as the zone-based ε -NTU method, have been described in detail in Naduvilakath-Mohammed *et al.* [26], so will not be repeated here. However, one key difference is worth mentioning as it pertains to the *Heat Recovery Model*, since

this requires a set of correlations for predicting the condensation heat transfer on the plate heat exchanger shown in Figure 5.3. For completeness, this has been added in Appendix A. Furthermore, the *Heat Recovery Model* will be used in the system design case studies to follow, and a key performance target parameter is the outlet temperature, T_{13} , of the single-phase water coolant (*i.e.* heat recovery fluid). Due to sensible heat transfer, T_{13} is sensitive to the total heat transfer rate from the refrigerant, as well as the mass flow rate, given that,

$$q_c = q_{\text{CPU}} + (W_{\text{pump}} + W_{\text{comp}}) = \dot{m}_{\text{hrf}} C_{p,\text{H}_2\text{O}} (T_{13} - T_{12}) \quad (5.11)$$

where $\dot{m}_{\text{hrf}} = \rho_{\text{H}_2\text{O}} \dot{V}_{\text{hrf}}$ is the mass flow rate of the heat recovery fluid for a volumetric flow rate of $\dot{V}_{\text{hrf}}$.

5.3.2.3. Expansion Device and Superheat

Assuming adiabatic conditions, the expansion process is isenthalpic such that the refrigerant expands from the condenser pressure to the evaporator pressure such that,

$$h_5 = h_6 \quad (5.12)$$

For the *Validation Model*, the throttling device is a simple capillary tube, and the degree of superheat ($\Delta T_{\text{sh}} = T_1 - T_7$) is determined empirically as,

$$\Delta T_{\text{sh}} = 12.85 + 0.0037S - 3.17P_e - 2.33P_c + 0.815P_e P_c \quad (\text{Validation Model}) \quad (5.13)$$

which fits the data with an adequate $R^2 = 0.83$. Although capillary tube throttling devices are simple and functional, they tend to maintain an unnecessarily high degree of refrigerant superheating, which is well known to be deleterious to overall system COP. Therefore, for the *Heat Recovery Model*, which will subsequently be used in system design case studies which includes energetic optimization, a Thermostatic Expansion Valve (TXV) type throttling device is implemented, since this maintains a relatively low and constant degree of superheat. To facilitate this, the TXV formulation used in Naduvilakath-Mohammed *et al.* [26] is adopted and

adapted for the present system, centered around 5°C, since the empirical correlation had already been developed as,

$$\Delta T_{sh} = 5.3 + 0.005T_e - 0.0082T_c + 0.000052S \quad (\text{Heat Recovery Model}) \quad (5.14)$$

5.3.3 Liquid-Cooling Loop

The liquid-cooling loop is modelled using the exactly the same approach as Naduvilakath-Mohammed *et al.* [26] and is reproduced here for completeness. The CPU cooling performance of the system is monitored by the die temperature, T_{die} . For a given CPU/TTV heat loading, q_{CPU} , this is predicted by,

$$T_{die} = T_9 + R_{th,clp}q_{CPU} \quad (5.15)$$

where T_9 is the inlet temperature of the coolant to the cold plate heat exchanger. $R_{th,clp}$ is the die-to-coolant thermal resistance, which includes the influence of the heat spreader, the layer of Thermal Interface Material (TIM), and the cold plate heat exchanger. A straight-forward correlation for the thermal resistance, which includes a term to account for changes in thermophysical properties with coolant temperature, was developed in Ref [26] as,

$$R_{th,clp} = 0.007894\dot{V}_{liq}^{-0.166} \left(\frac{\mu_{liq}}{k_{liq}} \right)^{-0.189} \quad (5.16)$$

where $\dot{V}_{liq}$ is the liquid coolant volumetric flow rate. Similarly, the empirical correlation for the pressure drop across the cold plate is adopted directly from Ref [26], as;

$$\Delta P_{liq,clp} = 2.45\dot{V}_{liq}^2 \left(\frac{\mu_{liq}}{\mu_{ref}} \right)^{-0.513} \quad (5.17)$$

5.3.4. System Energetics: Heat Recovery Model

As just discussed, from a CPU cooling standpoint, the described system acts as a refrigeration system to maintain acceptable CPU die temperatures for a specified CPU heat flux, with T_{die} serving as the primary performance indicator. From a refrigeration perspective, the Coefficient of Performance is related to the cooling load in the evaporator, q_e , and the total Power Consumption, PC, of the

system. Concerning the *Heat Recovery Model*, PC is determined by the sum of the power consumed by compressor, P_{comp} , together with the hydraulic power associated with the coolant flow in both the primary and heat recovery loops, W_{pump} and $W_{\text{hyd,hrf}}$ respectively. Thus, the net fluid system power consumption can be expressed as,

$$PC = P_{\text{comp}} + W_{\text{pump}} + W_{\text{hyd,hrf}} \quad (5.18)$$

The compressor electrical power P_{comp} is discussed in Section 3.2.1, as was the hydraulic power required to circulate the coolant through primary coolant loop. Similar to the latter, the hydraulic power penalty associated with the heat recovery fluid flow in the condenser is,

$$W_{\text{hyd,hrf}} = \dot{V}_{\text{hrf}} \Delta P_{\text{hrf}} \quad (5.19)$$

where ΔP_{hrf} is the pressure drop across the condenser heat exchanger. Defining the refrigeration Coefficient of performance is thus,

$$\text{COP} = \frac{q_e}{PC} = \frac{q_{\text{CPU}} + W_{\text{pump}}}{P_{\text{comp}} + W_{\text{pump}} + W_{\text{hyd,hrf}}} \quad (5.20)$$

Also of interest in this study is the vapour-compression system acting as a heat pump, since a key aim of the *Heat Recovery Model* is to investigate recovering the thermal energy at high temperature. In this case, the Coefficient of Performance is defined in terms of the total heat transfer rate to the heat recovery fluid in the condenser, q_c such that the heat pump COP is,

$$\text{COP}_{\text{HP}} = \frac{q_c}{PC} = \frac{q_e + W_{\text{comp}}}{P_{\text{comp}} + W_{\text{pump}} + W_{\text{hyd,hrf}}} \quad (5.21)$$

5.3.5. Model Validation

Experiments were performed over the following range of input parameters: CPU heat flux between 75 - 135 W/cm², compressor speed between 2174 - 4500 rpm, air speeds between 0.9 - 1.1 m/s, and primary coolant liquid flow rates between 1 - 6 L/min. Over the same range, the *Validation Model* was implemented to predict key state and performance parameters. Figure 5.5 compares the predicted and

measured system temperatures, pressures and COP, showing acceptable agreement over the full range of operating conditions. Notably, the data points for all monitored comparison metrics fall within the $\pm 5\%$ deviation bands. To quantify the level of agreement, the Mean Relative Deviation (MRD) of compared parameter M is calculated as,

$$\text{MRD} = \frac{1}{n} \sum |(M_{\text{exp}} - M_{\text{mod}})/M_{\text{exp}}| \times 100\% \quad (5.22)$$

where n is the number of test conditions.




Figure 5.5. Comparison of predicted and measured system (a) temperatures, (b) condenser and evaporator pressures and pressure ratio, and (c) Coefficient of Performance over a wide range of heat source powers, compressor speeds, air speeds, and liquid coolant flow rates. Predicted results obtained using the *Validation Model*.

The calculated MRD results are summarized in Table 5.3 and emphasize the acceptable agreement between the *Validation Model* predictions and experimental data. The good agreement between the *Validation Model* and the experimental data supports the modelling strategy used. Although no experimental data was produced for direct comparison with the *Heat Recovery Model*, it is assumed here that the addition of the plate heat exchanger condenser and TXV throttling device are minor enough that their integration with the DOE methodology for optimization and parameter sensitivity analysis is justified.

Table 5.3. Mean Absolute Relative Deviation (MARD) between the Validation Model predicted temperatures, pressures and system COP and experimentally determined values, over a wide range of operating conditions.

	P_e	P_c	T_1	T_2	T_5	T_6	T_9	T_{die}	COP
MRD (%)	0.5	0.7	0.2	0.15	0.15	0.18	0.19	0.17	0.17
MARD (%)	2.23	0.73	0.96	0.75	0.8	0.92	0.87	0.77	2.99

5.4. Design of Experiments

Following the same methodology developed in Naduvilakath-Mohammed *et al.* [26], a Design of Experiments (DOE) procedure was carried out using the Minitab commercial software. Using predicted results from the *Heat Recovery Model*, the overall DOE strategy follows the same framework established in Ref [26], but it is extended here to incorporate the additional thermodynamics associated with the heat recovery application.

The DOE procedure requires two key stages, an initial fractional factorial design to screen the most influential factors, followed by a central composite design process to capture curvature and interactions between variables within the defined operating envelope. Training of the DOE model is performed with seven factors: the CPU heat flux (q''), the primary-loop coolant flow rate (Q_{liq}), the compressor speed (S), the evaporator length (L_e), the refrigerant-liquid condenser length (L_c), the heat recovery liquid flow rate (Q_{hrf}), and the heat recovery liquid inlet temperature (T_{12}). The selected factor ranges are presented in Table 5.4.

Table 5.4. Summary of design factors and response variables used in the DOE model.

Category	Parameter	Unit	Range	Levels	Role/Remarks
Factors	Heat Flux, q''	W/cm ²	86 – 160	5	Disturbance / Design Load
	Liquid Flow Rate, Q_{liq}	L/min	0.1 – 6.0	6	Control / Fluid Mover
	Compressor Speed, S	rpm	2174 – 4500	4	Control / Fluid Mover
	HRF Flow Rate, Q_{hrf}	L/min	0.1 – 6.0	6	Control / Fluid Mover
	HRF Inlet Temperature, T_{12}	°C	60 – 70	2	Control / Fluid Mover
	Evaporator Length, L_e	cm	19.2 – 115.2	4	Control / Geometry
	Condenser Length, L_c	cm	19.2 – 115.2	4	Control / Geometry
Responses	Die Temperature, T_{die}	°C	—	—	Thermal Control Target
	HRF Outlet Temperature, T_{13}	°C	—	—	Thermal Control Target
	Power Consumption, PC	W	—	—	System Electrical Load

For the combined refrigeration - heat pump application of the Vapour-Compression system in this investigation, the DOE model is then utilized to predict three primary system responses: the CPU die temperature (T_{die}), the heat recovery fluid outlet temperature (T_{13}), and the total system power consumption (PC), as listed in Table 5.4. The system Coefficient of Performance (COP) is derived directly from PC, and is therefore not treated as a separate DOE response.

As will become clear in the following sections, implementing the *Heat Recovery Model* within a DOE framework not only serves as an effective means for analysing and understanding multiple response sensitivities to multiple input factors, but it is also useful as a design tool which can identify operating and design conditions that optimise system performance for given performance targets.

5.5. Results and Discussions

5.5.1. Parameter Sensitivity Analysis

A parameter sensitivity analysis using the DOE method is performed to examine how variations in the input factors influence the selected system responses (Table 5.4). Figure 5.6 presents the results and reveals that, considering all three responses, the CPU heat flux, q'' , is the most influential factor overall. Apart from it

being the obvious factor which most affects an increase in die temperature, T_{die} , (+40%), it also strongly influences the heat recovery fluid outlet temperature T_{13} (+23%), since an increased CPU power results in higher heat transfer rate to the heat recovery fluid in the condenser heat exchanger. It is noteworthy that the influence of CPU heat flux on power consumption, PC, is high (43%) since, for fixed compressor speed, a higher cooling load requires higher compressor power to accommodate the higher refrigerant mass flow rate [25].




Figure 5.6. Local sensitivity of system responses to operating and design factors.

In terms of the fluid movers, both PC and T_{die} are strongly influenced by the compressor speed, S . Its dominant influence on both energy and cooling performance is highlighted in the response surface plots in Figure 5.7(a) and 5.7(c) respectively, with trends consistent with those of Naduvilakath-Mohammed *et al.* [26]. The compressor speed is thus a key factor in balancing the trade-off between achieving cooling targets and maximizing system COP. In the context of attaining high heat recovery fluid outlet temperature, T_{13} , Figure 5.6 together with the response surface plot in Figure 5.7(b) indicate that increasing the compressor speed has a positive, though moderate influence (+11%). Increasing the compressor speed increases the compressor power, and thus the overall heat rejection in the condenser. However, compared with the evaporator cooling load,

q_e , the relative magnitude of compressor power increase with S is generally several times smaller, which diminishes the overall sensitivity of T_{13} .




Figure 5.7. Response surface plots for (a) T_{die} versus Q_{liq} & S , (b) T_{13} versus Q_{liq} & S , (c) PC versus Q_{liq} & S , (d) T_{die} versus L_e & L_c , (e) T_{13} versus L_e & L_c , and (f) PC versus L_e & L_c . For plots (a)-(c), q'' , L_e , L_c , Q_{hrf} & T_{12} , and for (d)-(f), the q'' , S , Q_{liq} , Q_{hrf} & T_{12} are fixed parameter settings.

In contrast to Naduvilakath-Mohammed *et al.* [26], the current VCR-assisted cooling with heat recovery system has two liquid-cooling loops, one for cooling the CPU and one for recovering the energy in the condenser. With regard to the heat recovery fluid flow rate factor, Q_{hrf} , the only response sensitivity of note is that of T_{13}

(−22%). This reasonably high sensitivity is expected since, for a given heat transfer rate in the condenser, the liquid temperature rise is directly proportional to the flow rate. Furthermore, Figure 5.7 shows that PC has negligibly small sensitivity to Q_{hrf} . Though increasing the heat recovery fluid flow rate adds some minor hydraulic power, this is offset by it also increasing the condenser heat exchanger effectiveness, which tends to lower the condenser pressure, the resulting pressure ratio, and thus the compressor power. Concerning the CPU/evaporator-side liquid coolant flow rate, Q_{liq} , Figure 5.6 indicates that the general conclusion from the Naduvilakath-Mohammed *et al.* [25,26] study still holds, in that CPU-side coolant control offers effective die temperature regulation with comparatively small energy consumption penalty. Figure 5.7 (a) and 5.7(c) demonstrate this more clearly, showing that a two-way trade-off between Q_{liq} and S can be balanced in such a way that the most energetically efficient approach to CPU cooling is to increase Q_{liq} , allowing S to be lowered. From the point of view of heat recovery, the advantage of using high CPU coolant flow rate is compounded to some degree. Figure 5.6 shows that T_{13} has a small, though positive sensitivity to Q_{liq} (+4.5%), which is further emphasized in Figure 5.7(b). For the range of flow rates considered, increasing Q_{liq} increases the hydraulic power in the CPU-side single-phase flow loop, though this is relatively minor in comparison to the impact of CPU power and compressor speed. Regardless, increasing the hydraulic power does increase the evaporator cooling load, q_e , and thus the total heat transfer rate in the condenser, q_c , which is beneficial in terms of the heat recovery temperature, T_{13} . Overall, one anticipates that design scenarios that involve demanding CPU and heat recovery thermal targets with minimum system power consumption will default to high Q_{liq} and low Q_{hrf} .

Consistent with Naduvilakath-Mohammed *et al.* [26], Figure 5.6 and Figure 5.7(d) show that increasing L_e strongly influences T_{die} . In fact, it is the second most influential factor to improve CPU cooling, with S being the most effective factor (Figure 5.6). Similar to Q_{liq} , Figure 5.6 and Figure 5.7(d) & (f) illustrate that the notable CPU cooling benefit is associated with a modest impact on PC (+4%, Figure 5.6), signifying that enlarging the evaporator also provides an energetically efficient means of CPU cooling, though in this case at the expense of overall system compactness [26]. In terms of influencing the heat recovery factor, T_{13} , Figure 5.6

and Figure 5.7(e) show a moderate positive influence with increased size of the evaporator (+6%, Figure 5.6). Similar to the discussion just above regarding Q_{liq} , increasing L_e increases the hydraulic power in the CPU-side single-phase loop, and thus the cooling load, q_e . The knock-on effect is a higher heat load in the condenser heat exchanger, which accounts for modest, yet positive, influence on T_{13} . Conversely, Figure 5.6 and Figure 5.7(d) demonstrate that increasing the condenser length, L_c , has very small influence on T_{die} (+1.8%, Figure 5.6), though is the only noteworthy factor that has an influence on reducing the overall power consumption (-4%, Figure 5.6). Increasing the condenser length increases the heat exchanger effectiveness, so for the same reasons though more pronounced than the effect of increasing Q_{hrf} discussed above, the knock-on effects are lower pressure ratio and compressor power. Although these are beneficial in the context of system energetics, Figure 5.6 and Figure 5.7(e) show that increasing the condenser length has a detrimental effect on T_{13} (-4%, Figure 5.6) due to the influence which tends to lower the net heat load in the condenser. Overall, the influence of the condenser on the response factors is small, and in the context as a design parameter, will likely be a fine-tuning parameter.

The final factor is the approach temperature of the heat recovery fluid to the condenser heat exchanger, T_{12} . For a given condenser heat load and liquid flow rate, Q_{hrf} , T_{13} increases by the same level as T_{12} , so the high sensitivity of T_{13} to T_{12} is unsurprising. However, for the range of T_{12} considered here, both the T_{die} and PC responses show very low sensitivity to T_{12} , implying that target heat recovery temperatures can be achieved with small penalties. However, the maximum outlet pressure is constrained to ~27 MPa for the compressor used in this study, which thus constrains T_{12} to an envelope where this maximum pressure is not exceeded.

5.5.2. Design Optimization Case Study

Considering the above discussion, the VCR-assisted CPU cooling and heat recovery system considered here has multiple design and control factors, and each one influences the three selected responses (T_{die} , T_{13} , and PC) differently, by both level of sensitivity and direction. In terms of the system design optimization case studies to follow, some factors have minor to moderate influence (*e.g.* L_e and L_c), and it is anticipated that these will be more akin to fine-tuning parameters when

balancing the trade-off to achieve thermal design targets with highest energetic efficiency. Other factors have obvious and significant response-specific influence, such as Q_{hrf} and T_{12} on the heat recovery fluid temperature, T_{13} . In terms of system control settings, it is likely that, similar to the Naduvilakath-Mohammed *et al.* [26] study, the key to achieving thermal targets with highest energetic efficiency will be the balancing of the compressor speed, S , and the CPU-loop coolant flow rate, Q_{liq} , especially considering its positive influence on both T_{die} and T_{13} . The balancing of the multiple factors with the inherent trade-offs as they influence each the CPU cooling, heat-recovery capability and energetic performance, is clearly challenging. However, as will be demonstrated below, the DOE-integrated model developed here streamlines this task, enabling rapid identification of control settings and design parameters that optimize the energy performance for specific cooling and heat recovery targets.

To demonstrate this capability, the DOE framework has been applied to selected case study situations. Table 5.5 summarises the three selected case studies, which span a range of design scenarios to not only demonstrate the efficacy of the developed optimization design approach, but also to allow some meaningful insight into the thermodynamics of the system. In these demonstration cases, the heat recovery fluid inlet temperature is fixed at $T_{12} = 70^{\circ}\text{C}$, as it was determined that this was sufficiently high to achieve the target heat recovery fluid outlet temperature of $T_{13} = 85^{\circ}\text{C}$, whilst not exceeding the compressor exit pressure constraint of 27 MPa. The design objectives for these case studies are as follows:

Case 0: For a selected high heat flux ($q'' = 86 \text{ W/cm}^2$), fixed evaporator and condenser sizes ($L_e = L_c = 38.4 \text{ cm}$), and fixed heat recovery fluid inlet temperatures ($T_{12} = 70^{\circ}\text{C}$), the objective is to achieve target CPU die and heat recovery fluid temperatures ($T_{die} = T_{13} = 85^{\circ}\text{C}$) whilst minimizing PC (maximizing COP). Here the fluid mover settings (Q_{liq} , Q_{hrf} , and S) are unconstrained. This is considered the baseline case.

Case 1: For a selected high heat flux ($q'' = 86 \text{ W/cm}^2$) and fixed heat recovery fluid inlet temperature ($T_{12} = 70^{\circ}\text{C}$), the objective is to achieve target CPU die and heat recovery fluid temperatures ($T_{die} = T_{13} = 85^{\circ}\text{C}$) whilst minimizing PC (maximizing COP). Here the fluid mover (Q_{liq} , Q_{hrf} , and S) and heat exchanger

sizes (L_e and L_c) are unconstrained. This case demonstrates how drop-in replacement heat exchangers can influence the CPU cooling, heat recovery and COP relative to the baseline.

Table 5.5. DOE optimization case study test matrix.

Case	q''	S	Q_{liq}	Q_{hrf}	T_{12}	L_e	L_c	T_{die}	T_{13}	PC	Optimization
0	HV	UC	UC	UC	HV	HV	HV	TG	TG	Min	SOO
1	HV	UC	UC	UC	HV	UC	UC	TG	TG	Min	SOO
2	HV	UC	UC	UC	HV	UC	UC	TG	TG	Min	SOO

HV=Hold Value, TG=Target Goal, UC=Unconstrained, Min=Minimize Goal, SOO=Single-Objective Optimization.

Case 2: For a selected maximum heat flux ($q'' = 156 \text{ W/cm}^2$) and fixed heat recovery fluid inlet temperature ($T_{12} = 70^\circ\text{C}$), the objective is to achieve the target CPU throttling temperature ($T_{die} = 100^\circ\text{C}$) and heat recovery fluid temperature ($T_{13} = 85^\circ\text{C}$) whilst minimizing PC (maximizing COP). As in Case 1, the fluid mover (Q_{liq} , Q_{hrf} , and S) and heat exchanger sizes (L_e and L_c) are unconstrained. This case demonstrates system cooling performance for challenging heat loads while maintaining heat recovery goals with optimized energy efficiency.

Table 5.6 summarises the control factors and system responses for the three optimisation demonstration case studies. The baseline Case 0 reveals some of the parameter sensitivity talking points discussed above. In this scenario, with the fixed heat exchanger sizes, only the fluid mover settings (Q_{liq} , Q_{hrf} , and S) are unconstrained to achieve the $T_{die} = T_{13} = 85^\circ\text{C}$ thermal targets with minimal overall system power consumption. In the previous studies of Naduvilakath-Mohammed *et al.* [25,26], which concentrated solely on CPU cooling and energetic performance, the consistent trend was to increase Q_{liq} to its maximum and reduce S until the CPU thermal target was achieved. This provided the best trade-off between CPU cooling and overall system energy consumption. Table 5.6 shows that the same result is found here, with Q_{liq} being at its maximum of 6 L/min for Case 0, despite the heat

recovery fluid temperature now being included as a targeted response parameter. As discussed, increasing Q_{liq} has a minor, though positive influence on increasing T_{die} . Table 5.6 also shows that the heat recovery fluid flow rate is quite low, at $Q_{hrf} = 0.39$ L/min. This is expected since, as discussed, T_{13} has high and negative sensitivity to Q_{hrf} with little influence on the T_{die} and PC responses (Figure 5.6). Overall, by simply controlling the Q_{liq} , Q_{hrf} , and S fluid mover settings, the CPU and heat recovery targets can be achieved with a reasonably high COP = 4.0.

Table 5.6. DOE optimization case study results.

Case	Objective	Factors							Responses			
		Heat Flux, q'' (W/cm ²)	Compressor Speed, S (rpm)	Liquid Flow Rate, Q_{liq} (L/min)	HRF Flow Rate, Q_{hrf} (L/min)	HRF Inlet Temperature, T_{12} (°C)	Length (cm)		Die Temperature T_{die} (°C)	HRF Outlet Temperature T_{13} (°C)	Power Consumption PC (W)	Coefficient of Performance COP (-)
							Evaporator L_e	Condenser L_c				
0	Baseline	86 [Hold Value]	2549	6	0.39	70 [Hold Value]	38.4 [Hold Value]	38.4 [Hold Value]	85 [Goal: Target]	85 [Goal: Target]	87.8 [Goal: Minimize]	4.0 [Goal: Maximize]
1	Energy Efficiency - Heat Recovery	86 [Hold Value]	2174	6	0.36	70 [Hold Value]	67.2	28.5	85 [Goal: Target]	85 [Goal: Target]	69.8 [Goal: Minimize]	5.1 [Goal: Maximize]
2	Extreme Heat Load - Heat Recovery	156 [Hold Value]	2501	6	0.72	70 [Hold Value]	115.2	115.2	100 [Goal: Target]	85 [Goal: Target]	221.5 [Goal: Minimize]	3.0 [Goal: Maximize]

Relative to the baseline Case 0, Table 5.6 shows that the energy-efficiency and heat-recovery optimised design Case 1 can attain identical CPU cooling and heat recovery thermal targets with ~28% improvement in COP (4.0 → 5.1). As the table indicates, the primary factors that differentiate Case 1 from Case 0 are the compressor speed, S, and the evaporator and condenser sizes, L_e and L_c respectively, remembering that the latter two are now unconstrained. Overall, the improved COP of Case 1 is due to a 21% reduction in the system power consumption, PC (87.8 W → 69.8 W), which is primarily due to the compressor dropping to its minimum allowable speed. In order to dissect how this design achieves the CPU and heat recovery thermal targets with an improvement in COP, the discussion will be broken-down below, with aid of Figure 5.8 - Figure 5.10:

- i. For Case 1, the CPU liquid coolant flow rate is identical to that of Case 0, $Q_{liq} = 6$ L/min, and for the same reasons. Thus, the double benefit of Q_{liq}

with T_{die} and T_{13} is again such that the maximum flow rate is predicted. Furthermore, with the same flow rate, the thermal resistance of the cold plate heat exchanger cooling the CPU is the same for both test cases.

- ii. The compressor speed has dropped from $S = 2946$ rpm to its minimum speed in the selected range, $S = 2174$ rpm. As Figure 5.10 shows, this results in a 31% drop in compressor power ($78.6 \text{ W} \rightarrow 54.6 \text{ W}$).




Figure 5.8. P-h diagram comparing Cases 0 & 1, with $q'' = 86 \text{ W/cm}^2$, $T_{die} = T_{13} = 85^\circ\text{C}$.

- iii. As discussed earlier and shown in Table 5.6, decreasing S has a sizable effect tending to increase T_{die} . Even with Q_{liq} at its maximum, S at its minimum setting is not sufficient to achieve the CPU thermal target. Thus, it is required that the size of the evaporator be increased by ~ 1.8 -fold, which as discussed, is an energetically efficient means of reducing T_{die} . Albeit a minor influence, it is again worth noting that increasing L_e tends to have a positive influence on the heat recovery temperature, T_{13} . Thus,

increasing L_e influences both thermal targets positively, with low penalty in terms of impact on PC (Figure 5.6).




Figure 5.9. Evaporator and condenser performance parameters (T_3 , T_6 , T_9 , T_{13} , ϵ) for demonstration Cases 0 and 1, for $q'' = 86 \text{ W/cm}^2$, $T_{\text{die}} = T_{13} = 85^\circ\text{C}$.

- iv. Examining the corresponding cycles on the P-h diagram, Figure 5.8 shows that the larger evaporator of Case 1 elevates the evaporator pressure, and thus the saturation temperature T_6 ($13.0^\circ\text{C} \rightarrow 16.3^\circ\text{C}$, Figure 5.9). However, to achieve the same target die temperature of $T_{\text{die}} = 85^\circ\text{C}$ under the same CPU heat load and with the same cold plate thermal resistance, the liquid coolant temperature exiting the evaporator, T_9 , must be the same for both test cases. Figure 5.9 shows that this is indeed true, where $T_9 \sim 20^\circ\text{C}$ for both. However, as Figure 5.9 also indicates, the larger evaporator of Case 1 achieves the equivalent T_9 with the elevated refrigerant temperature due to the increased heat exchanger effectiveness, ϵ_{evap} ($0.69 \rightarrow 0.79$) [26].
- v. With regard to the condenser, Table 5.6 indicates that the Case 1 condenser heat exchanger is $\sim 20\%$ smaller than that of the baseline Case

0. The main reason for the drop in L_c is the net reduction in the compressor work, which subsequently causes a net decrease in the heat transfer rate in the condenser (434 W $\rightarrow$ 414 W). The reduction of L_c increases the condenser temperature (75.4 °C $\rightarrow$ 79.1°C, Figure 5.9) with only a small increase in condenser pressure (23.4 MPa $\rightarrow$ 25.8 MPa, Figure 5.8). The net influence is an increase in the driving temperature differential and thus local heat flux between the refrigerant and heat recovery coolant, resulting in only a minor (4%) drop in condenser heat exchanger effectiveness (0.52 $\rightarrow$ 0.5, Figure 5.9), despite the much larger change in the heat exchanger size. Secondly, the heat recovery fluid flow rate in the condenser, Q_{hrf} , drops marginally to ensure the lower sensible heat transfer rate raises its temperature to the required T_{13} target.




Figure 5.10. Power consumption budget breakdowns for demonstration Cases 0 and 1, for $q'' = 86 \text{ W/cm}^2$, $T_{die} = 85^\circ\text{C}$ and $T_{13} = 85^\circ\text{C}$.

Figure 5.10 shows a summary of the power consumption breakdown for each Cases 0 and 1, along with the associated COP. Comparatively, the most

prominent feature is the substantial reduction in compressor power ($78.6 \text{ W} \rightarrow 54.69 \text{ W}$), which is large enough to more than offset the increase in the hydraulic power of the CPU-side liquid coolant ($8.4 \text{ W} \rightarrow 14.7 \text{ W}$) as a result of the larger evaporator. It is also noteworthy that, as discussed above, the heat recovery liquid flow rate tends to be low to ensure that the required temperature lift ($70^\circ\text{C} \rightarrow 85^\circ\text{C}$) is achieved for a given sensible heating load. This plays positively to the overall system power consumption, since the hydraulic penalty associated with heat recovery liquid flow is negligibly small ($< 1 \text{ W}$).

Overall, the above comparison of test Cases 0 and 1 demonstrate that, for a given CPU heat load, appropriate heat exchanger sizing coupled with optimal control of the fluid-movers can simultaneously achieve both CPU cooling and heat recovery thermal targets whilst reducing the total power consumption, thus resulting in an improvement in system COP.

A final case is now considered whereby the CPU is at a very high heat flux of $q'' = 156 \text{ W/cm}^2$ (624 W), with the die temperature allowed to reach the approximate throttling temperature ($T_{\text{die}} = 100^\circ\text{C}$). These are identical to the extreme heat load test case considered in the Naduvilakath-Mohammed *et al.* [26] study. Here however, the added response target of $T_{13} = 85^\circ\text{C}$ is included, with the overarching objective still being to maximize the system COP. This corresponds to Case 2 in Table 5.6, and was selected to represent a particularly demanding CPU cooling scenario, such that the system must simultaneously manage very high thermal loads whilst delivering effective heat recovery. In the analysis that follows, this case is compared directly with the Case 1, as both allow unconstrained adjustment of the evaporator and condenser heat exchanger sizes (Table 5.5).

Relative to Case 1, Table 5.6 shows that the extreme heat load with heat recovery optimised design, Case 2, can achieve the more demanding CPU cooling with a 41% drop in COP ($5.1 \rightarrow 3.0$). As the table shows, the main factors that distinguish between Cases 1 and Case 2 are the compressor speed, S , the heat exchanger sizes, L_e & L_c , and the heat recovery fluid flow rate, Q_{hrf} . Overall, the decreased COP is due to 3-fold increase of the system total power consumption, PC ($69.8 \text{ W} \rightarrow 221.5 \text{ W}$), which is primarily due to the higher cooling load and increased compressor speed, which Table 5.6 shows have high influence on PC.

Similar to the discussion format above, the key aspects which differentiate Case 1 and Case 2 will be broken-down below with aid of Figure 5.11 and Figure 5.12.

- (1) For Case 2, the CPU liquid coolant flow rate is again at its maximum of 6 L/min. This minimizes the thermal resistance of the cold plate heat exchanger cooling the CPU with no added penalty in associated hydraulic power.




Figure 5.11. P-h diagram comparing Cases 1 & 2, with $T_{13} = 85^{\circ}\text{C}$

- (2) The compressor speed has increased moderately (2174 rpm $\rightarrow$ 2526 rpm). Despite a 15% increase in speed, Figure 5.12 shows that the compressor power has over tripled, (55 W $\rightarrow$ 192 W). Referring to the P-h diagram in Figure 5.11, the enthalpy difference across the evaporator is nearly the same for both cases. As such, the higher evaporator heat load of Case 2 generates a higher refrigerant mass flow rate (0.0031 kg/s $\rightarrow$ 0.0057 kg/s), which in turn demands disproportionately large increase in compressor power for a modest increase in compressor speed (Figure 5.6). It is worth noting that increasing the compressor speed generally tends to decrease the suction-side pressure, though referring to Figure 5.11, this is not the case here.

- (3) The evaporator size increases by ~71% to the largest allowable, for reasons already discussed. This accounts for the calculated increase the evaporator effectiveness ($0.79 \rightarrow 0.86$), causing the evaporator pressure to increase (Figure 5.11), despite the increased compressor speed. Overall, this moderates the required compressor power at the substantially higher evaporator loading by reducing the pressure ratio ($5.1 \rightarrow 4.8$), with a low penalty in terms of additional evaporator hydraulic power ($14.7 \text{ W} \rightarrow 25.2 \text{ W}$, Figure 5.12).




Figure 5.12. Power consumption budget breakdowns for demonstration Cases 1 and 2, for $T_{13} = 85^\circ\text{C}$

- (4) The increase in the evaporator and compressor power of Case 2 compared with Case 1 causes the heat transfer rate in the condenser to increase by more than 2-fold ($414 \text{ W} \rightarrow 840 \text{ W}$). However, Figure 5.11 shows that the condenser temperatures, T_3 , are in fact nearly the same ($79.1^\circ\text{C} \rightarrow 80.5^\circ\text{C}$), and this is due to the larger condenser size, which has increased 4-fold for Case 2. This, together with improved convective heat transfer coefficients, causes the condenser heat exchanger effectiveness, ϵ_{cond} , to increase substantially ($0.5 \rightarrow 0.8$). The overall effect is that the higher heat load can be dissipated at the

same condenser temperature, which subsequently moderates the condenser pressure and associated compressor work (Table 5.6). Secondly, the heat recovery fluid flow rate, Q_{hrf} , nearly doubles to accommodate the higher condenser heat transfer rate whilst achieving the $T_{13} = 85^{\circ}\text{C}$ target.




Figure 5.13. Heat pump Coefficient of Performance for demonstration Cases 0, 1 and 2.

Figure 5.12 gives a summary plot of the power consumption breakdown and COPs for Cases 1 and 2. Although the hydraulic power associated with the liquid-side coolants in both the evaporator and condenser heat exchangers increase, these pale in comparison with the compressor power increase. Despite this, it is worth now mentioning that, although the discussion has focussed on the vapour-compression system attaining high quality (temperature) heat, the quantity of heat recovered is also important, particularly for applications such as DAC where both the heat supply temperature and thermal duty are relevant. Thus, the observed increase in overall system power consumption is on one hand a negative, since it tends to reduce the refrigeration-based COP, but on the other hand is a positive, since it increases the quantity of heat energy available for reuse in downstream heat recovery applications. Thus, when considered in the context of a heat pump, both

the heat recovery temperature and COP_{HP} (Eq. 5.19) can both be considered relevant. Figure 5.13 plots COP_{HP} for each demonstration Cases 0, 1 and 2, remembering that for each case, the main heat pump function considered thus far is the production of hot water at $T_{13} = 85^\circ\text{C}$. As the figure indicates, when viewed from the perspective as a heat pump, the energetic performance is reasonably good, with $3.8 \leq \text{COP}_{\text{HP}} \leq 6$.

5.6. Conclusions

This work investigated a miniature vapour-compression assisted liquid system for cooling high-power CPUs whilst producing high temperature ($> 80^\circ\text{C}$) hot water for potential reuse of recovered thermal energy in Direct Air Capture (DAC) systems. With R-134a as the refrigerant and a high pressure ratio compressor, experiments were performed on an air-cooled vapour-compression system. The collected data was subsequently used to validate a full system mathematical model. The validated model was then augmented to incorporate a refrigerant-to-water condenser heat exchanger to better represent a heat recovery suitable design. The model was then used in conjunction with a Design of Experiments (DOE) methodology to investigate system response sensitivities of CPU cooling (T_{die}), heat recovery (T_{13}), and energetic (PC/COP) performance to input parameters. The DOE was then used to perform system-level optimization on three case studies to demonstrate the utility of the proposed design methodology. The main findings are as follows:

- (1) For a given CPU heat loading, the compressor speed is the dominant parameter across all selected responses, T_{die} , T_{13} and PC. Therefore, reaching challenging CPU and heat recovery thermal targets can come at the expense of high compressor power consumption, which is detrimental to system COP.
- (2) Compressor power requirements can be reduced through system-level optimisation aimed at maximising overall COP.
- (3) The analysis shows that appropriate sizing of the vapour-compression refrigeration heat exchangers, combined with effective balancing of the fluid-

mover operating conditions, can result in a substantial improvement in COP for a given CPU and heat recovery thermal targets.

- (4) With respect to managing the trade-off between cooling performance and energy consumption, whilst achieving target heat recovery temperatures, all of the design factors adjusted in such a way as to mitigate excessive compressor power consumption. It was found that increasing the liquid coolant flow rate and/or evaporator size has a disproportionately strong positive impact on CPU cooling effectiveness and heat recovery temperature, relative to their associated energy penalty. This enables CPU and heat recovery thermal targets to be met while maximising COP. The condenser heat exchanger size and heat recovery liquid flow rate scaled with condenser heat load in order to maintain a relatively constant condenser pressure for the target outlet coolant temperature of $T_{13} = 85^{\circ}\text{C}$.
- (5) From a CPU cooling perspective, where the system acts to refrigerate the liquid coolant, optimum COPs in the range of 3.0 - 5.1 were predicted, which are reasonably good considering the demanding thermal targets of the selected case studies. From a heat recovery perspective, where the system is a heat pump used to achieve heat recovery temperature of 85°C , the associated COP_{HP} values ranged between 3.8 - 5.9, which is evidence to suggest that the proposed technological approach to CPU cooling with high-temperature heat recovery is thermodynamically viable.
- (6) The selected case studies clearly demonstrate that the proposed mathematical model-DOE methodology is a viable strategy for designing and operating miniature vapour-compression assisted liquid-cooling with heat recovery systems for optimal performance.

While the recovered heat temperatures achieved in this study are comparable to those reported for certain DAC sorbent regeneration processes, the present work considers temperature as the primary heat recovery metric and does not evaluate the thermal duty requirements of practical DAC systems. Consequently, the results should be interpreted as demonstrating the capability of the miniature vapour-compression system to achieve DAC-relevant heat recovery

temperatures rather than establishing the feasibility of directly supporting commercial DAC operation. Assessment of the required heat supply rates, CO₂ capture capacities, and system-level integration between data-centre cooling systems and DAC processes remains an important area for future research.

A potential shortcoming of the present work relates to the use of R-134a. This is a high GWP refrigerant which is currently being phased-out, and is indeed not the ideal refrigerant for higher temperature heat recovery applications. Future work will explore the influence of different refrigerants on system thermal and energetic performance. Due to the very low charge volumes associated with this miniature system, even R-290 is a potential candidate. It is a natural refrigerant with very low GWP, tends to outperform synthetic refrigerants, and is currently widely deployed in domestic heat pump applications.

5.7. References

- [1] E. Masanet, A. Shehabi, N. Lei, S. Smith, J. Koomey, Recalibrating global data center energy-use estimates, *Science* 367 (2020) 984–986.
- [2] A. Abdilla, S.P. Borg, J. Licari, Relating measured PUE to the cooling strategy and operating conditions through a review of a number of Maltese data centres, *Energy Reports* 13 (2025) 2612–2623.
- [3] D. Donnellan, A. Lawrence, D. Bizo, P. Judge, J. O'Brien, J. Davis, M. Smolaks, J. Williams-George, R. Weinschenk, The Uptime Institute Global Data Center Survey, (2024).
- [4] M. Azarifar, M. Arik, J.-Y. Chang, Liquid cooling of data centers: A necessity facing challenges, *Applied Thermal Engineering* 247 (2024) 123112.
- [5] A. Carbó, E. Oró, J. Salom, M. Canuto, M. Macías, J. Guitart, Experimental and numerical analysis for potential heat reuse in liquid cooled data centres, *Energy Conversion and Management* 112 (2016) 135–145.
- [6] Directive (EU) 2023/1791 of the European Parliament and of the Council of 13 September 2023 on energy efficiency and amending Regulation (EU) 2023/955 (recast) (Text with EEA relevance), 2023.

- [7] Data Centre Energy and Sustainability Performance Reporting Obligations, Department of Climate, Energy and the Environment, Ireland, February, 2025
- [8] Deirde Nagle, Jay Sattin, Eoin Cassidy, New Data Centre Sustainability Reporting Obligations Introduced, Mason Hayes Curran, April, 2024.
- [9] M. Fasihi, O. Efimova, C. Breyer, Techno-economic assessment of CO₂ direct air capture plants, *Journal of Cleaner Production* 224 (2019) 957–980.
- [10] C.D. Díaz-Marín, Z.J. Berquist, Flipping the switch: carbon-negative and water-positive data centers through waste heat utilization, *Energy Environ. Sci.* 18 (2025) 8403–8413.
- [11] H.M. Ljungqvist, M. Risberg, A. Toffolo, M. Vesterlund, A realistic view on heat reuse from direct free air-cooled data centres, *Energy Conversion and Management: X* 20 (2023) 100473.
- [12] K. Ebrahimi, G.F. Jones, A.S. Fleischer, A review of data center cooling technology, operating conditions and the corresponding low-grade waste heat recovery opportunities, *Renewable and Sustainable Energy Reviews* 31 (2014) 622–638.
- [13] M. Wahlroos, M. Pärssinen, J. Manner, S. Syri, Utilizing data center waste heat in district heating – Impacts on energy efficiency and prospects for low-temperature district heating networks, *Energy* 140 (2017) 1228–1238.
- [14] E. Oró, P. Taddeo, J. Salom, Waste heat recovery from urban air cooled data centres to increase energy efficiency of district heating networks, *Sustainable Cities and Society* 45 (2019) 522–542.
- [15] W.-X. Chu, C.-C. Wang, A review on airflow management in data centers, *Applied Energy* 240 (2019) 84–119.
- [16] P. Kasten, S. Zimmermann, M.K. Tiwari, B. Michel, D. Poulikakos, Hot water cooled heat sinks for efficient data center cooling: Towards electronic cooling

- with high exergetic utility, *Frontiers in Heat and Mass Transfer* (2010) 1(2), 1-10.
- [17] S. Zimmermann, I. Meijer, M.K. Tiwari, S. Paredes, B. Michel, D. Poulikakos, Aquasar: A hot water cooled data center with direct energy reuse, *Energy* 43 (2012) 237–245.
- [18] N. Meyer, M. Ries, S. Solbrig, T. Wettig, *iDataCool: HPC with Hot-Water Cooling and Energy Reuse*, 2013.
- [19] R. Tu, L. Wang, L. Liu, X. Yang, Performance analyses and optimization design of a heat recovery system used between data centers and centralized heating systems, *Building Simulation* 18, 2559–2576 (2025).
- [20] L. Zhang, Z. Zhao, B. Chen, M. Zhao, Y. Chen, Zero-Carbon Development in Data Centers Using Waste Heat Recovery Technology: A Systematic Review, *Sustainability* 17 (2025) 10101.
- [21] T. Wilde, M. Ott, A. Auweter, I. Meijer, P. Ruch, M. Hilger, S. Kuhnert, H. Huber, CoolLMUC-2: A supercomputing cluster with heat recovery for adsorption cooling, 2017, 33rd Thermal Measurement, Modeling & Management Symposium (SEMI-THERM) (2017): 115-121.
- [22] M. Baghsheikhi, M. Mohammadi, Experimental investigation of the vapor-compression cooling system in a data center: energy and exergy analysis, *J Therm Anal Calorim* 148 (2023) 9079–9097.
- [23] J.C. Silva Romero, J. Belman-Flores, S. Aceves, A Review of Small-Scale Vapor Compression Refrigeration Technologies, *Applied Sciences* 14 (2024) 3069.
- [24] H. Zhou, J. Li, Development and analysis of a simple structured and economic miniature vapor compression refrigerator for cooling electronics in harsh environment, *Applied Thermal Engineering* 223 (2023) 120047.
- [25] F.M. Naduvilakath-Mohammed, M. Lebon, G. Byrne, A.J. Robinson, An experimental investigation of vapor compression refrigeration cooling and

- energy performance for CPU thermal management, *Case Studies in Thermal Engineering* 54 (2024) 103963.
- [26] F.M. Naduvilakath-Mohammed, S.S. Shojaei Zadeh, A.J. Robinson, Modelling and Optimization of Vapour-Compression Refrigerated Liquid Cooling of CPUs, *Case Studies in Thermal Engineering* (2026), submitted.
- [27] S. Mancin, C. Zilio, G. Righetti, L. Rossetto, Mini Vapor Cycle System for high density electronic cooling applications, *International Journal of Refrigeration* 36 (2013) 1191–1202.
- [28] S. Trutassanawin, E.A. Groll, S.V. Garimella, L. Cremaschi, Experimental Investigation of a Miniature-Scale Refrigeration System for Electronics Cooling, *IEEE Transactions on Components and Packaging Technologies* 29 (2006) 678–687.
- [29] Z. Wu, R. Du, Design and experimental study of a miniature vapor compression refrigeration system for electronics cooling, *Applied Thermal Engineering* 31 (2011) 385–390.
- [30] F.M. Naduvilakath-Mohammed, R. Jenkins, G. Byrne, A.J. Robinson, Closed loop liquid cooling of high-powered CPUs: A case study on cooling performance and energy optimization, *Case Studies in Thermal Engineering* 50 (2023) 103472.
- [31] J.R. Taylor, *An introduction to error analysis*, University Science Books, Sausalito, Calif., 1997.
- [32] J. Marchante-Avellaneda, J.M. Corberan, E. Navarro-Peris, S.S. Shrestha, A critical analysis of the AHRI polynomials for scroll compressor characterization, *Applied Thermal Engineering* 219 (2023) 119432.
- [33] J. Marchante-Avellaneda, E. Navarro-Peris, J.M. Corberan, S.S. Shrestha, Analysis of map-based models for reciprocating compressors and optimum selection of rating points, *International Journal of Refrigeration* 153 (2023) 168–183.

- [34] L.E. Olmedo, V. Mounier, L.C. Mendoza, J. Schiffmann, Dimensionless correlations and performance maps of scroll expanders for micro-scale Organic Rankine Cycles, *Energy* 156 (2018) 520–533.
- [35] E. Navarro, E. Granryd, J.F. Urchueguía, J.M. Corberán, A phenomenological model for analyzing reciprocating compressors, *International Journal of Refrigeration* 30 (2007) 1254–1265.
- [36] L. Dardenne, E. Fraccari, A. Maggioni, L. Molinaroli, L. Proserpio, E. Winandy, Semi-empirical modelling of a variable speed scroll compressor with vapour injection, *International Journal of Refrigeration* 54 (2015) 76–87.
- [37] M.-E. Duprez, E. Dumont, M. Frère, Modelling of reciprocating and scroll compressors, *International Journal of Refrigeration* 30 (2007) 873–886.
- [38] R. Ossorio, J. Marchante-Avellaneda, E. Navarro-Peris, Development of compact empirical models for variable-speed compressors for the prediction of energy consumption, mass flow and discharge temperature, *Applied Thermal Engineering* 244 (2024) 122666.
- [39] H. Kumar, *The Plate Heat Exchanger: Construction and Design*, Institute of Chemical Engineering Symposium Series No. 86 (1984) 1275–1288.
- [40] T.M. Adams, M.F. Dowling, S.I. Abdel-Khalik, S.M. Jeter, Applicability of traditional turbulent single-phase forced convection correlations to non-circular microchannels, *International Journal of Heat and Mass Transfer* 42 (1999) 4411–4415.
- [41] Y.-Y. Yan, H.-C. Lio, T.-F. Lin, Condensation heat transfer and pressure drop of refrigerant R-134a in a plate heat exchanger, *International Journal of Heat and Mass Transfer* 42 (1999) 993–1006.
- [42] W.W. Akers, H.A. Deans, O.K. Crosser, CONDENSING HEAT TRANSFER WITHIN HORIZONTAL TUBES, *Chem. Eng. Progr. Vol: 55, Symposium Ser. No. 29* (1958).

- [43] Y.A. Çengel, Heat and Mass Transfer: A Practical Approach, McGraw-Hill, 2007.
- [44] T.L. Bergman, F.P. Incropera, A.S. Lavine, D.P. DeWitt, Introduction to Heat Transfer, John Wiley & Sons, 2011.

Chapter 6

Conclusions and Future Work

6.1. Conclusions

This research systematically identified and addressed critical gaps in the understanding of advanced CPU thermal management technologies, particularly in relation to the growing limitations of air-cooling and single-phase liquid cooling in high power computing environments. Through the hypotheses and results developed across four integrated studies, the research provides new insight into the complex interplay between heat transfer processes, refrigeration-assisted cooling, fluid-mover penalties, system-level optimisation, and heat recovery potential. Collectively, these contributions establish a structured methodology for designing next generation cooling systems that deliver both enhanced thermal performance and high-temperature heat recovery. This chapter consolidates the key conclusions, linking them to the original gaps in knowledge and assessing the validity of the proposed hypotheses.

The first part of this research characterised the performance envelope of conventional fan–fin air cooling and hybrid liquid-air systems, establishing a robust baseline against which all later investigations into the cooling performance of vapour-compression assisted cooling is evaluated. Chapter 2 demonstrated that Hybrid Liquid-Air (HLA) cooling delivers substantially lower source-to-sink thermal resistance than air cooling and supports higher CPU heat loads while maintaining safe operating temperatures. The study further showed that coordinated control of pump and fan speeds can meaningfully reduce overall fluid-mover energy consumption, illustrating that non-refrigerated liquid cooling provides measurable performance advantages with acceptably low power consumptions. However, it was demonstrated that a key limitation to cooling performance was the source-to-sink thermal resistance being restrained to the ambient heat sink temperature.

Building on this baseline, Chapter 3 explored whether compact Vapour-Compression Refrigeration (VCR) could deliver sub-ambient coolant temperatures to further expand CPU thermal envelopes. The experimental results validated the

hypothesis that miniature VCR systems can reliably achieve sub-ambient conditions across a wide range of heat fluxes, significantly reducing CPU die temperatures relative to the more conventional HLA and fan-fin systems. While compressor operation introduced additional fluid-mover energy consumption penalties, the study demonstrated that careful selection of compressor speed, coolant flow rate, and condenser airflow enables favourable Coefficients of Performance (COP), up to $\text{COP} = 4.7$. This work provided the first rigorous, side-by-side comparison of all three cooling solutions within a unified experimental framework.

The third stage of the research addressed the need for a predictive, system-level mathematical modelling framework capable of capturing the nonlinear interactions between refrigerant thermodynamics, heat exchanger performance, and the primary liquid loop. A comprehensive steady-state mathematical model of the full VCR-assisted liquid-cooling system was developed, against experiments, then integrated with a Design of Experiments (DOE) methodology. The model predictions confirmed the hypothesis that compressor speed overwhelmingly governs cooling performance and energy consumption. At the same time, the results showed that increasing coolant flow rate or evaporator size yields disproportionately favourable cooling improvements relative to the incremental fluid-mover penalties introduced. This model-driven optimisation methodology developed demonstrated a robust engineering tool for analysing, designing, and tuning miniature VCR-assisted liquid-cooling systems for electronics thermal management.

The final research question investigated was whether miniature vapour-compression systems could be extended to support dual-purpose operation: maintaining safe CPU operation while producing the high-temperature heat recovery stream ($\geq 80^\circ\text{C}$) required for Direct Air Capture (DAC). To assess this, the previously developed modelling framework and methodology was expanded to incorporate a high pressure ratio compressor and water-cooled condenser, making it suitable for generating sub-ambient CPU-side coolant with elevated heat recovery temperature coolant. DOE-based optimisation validated the hypothesis that coordinated control of fluid-movers, and heat exchanger sizing allows simultaneous satisfaction of CPU cooling and DAC heat recovery temperature targets with

maximized system COP. Heat pump COPs between 3.8 and 5.9 demonstrated that energetically efficient operation remains attainable even with demanding CPU and heat recovery thermal targets. This marks the first demonstration that miniature vapour-compression assisted liquid-cooling systems can function concurrently as refrigeration units and high temperature heat pumps, offering a novel pathway for data centre decarbonisation.

Although the recovered heat temperatures achieved in this study are comparable to those reported for certain DAC regeneration processes, the present work considers temperature as the primary heat recovery metric and does not evaluate the thermal duty requirements of practical DAC systems. Consequently, the results should be interpreted as demonstrating DAC-relevant heat recovery temperatures rather than the feasibility of direct DAC integration.

Overall, this thesis delivers new theoretical insight, experimental validation, and modelling strategies for hybrid liquid cooling, VCR-assisted refrigeration, and high-temperature heat recovery within the broader field of electronics thermal management. By establishing a documented performance baseline, quantifying the benefits and penalties of VCR-assisted cooling, developing a validated simulation-optimisation methodology, and demonstrating dual-purpose cooling - heat recovery operation, the research provides a comprehensive foundation for the next generation of high-performance, energy-efficient, and sustainability-aligned cooling solutions. Through the integrated consideration of heat transfer behaviour, fluid-mover dynamics, and optimisation principles, this work positions itself as a significant contribution to the evolving landscape of data centre thermal engineering.

6.2. Future Work

This thesis has demonstrated the feasibility and limitations of hybrid VCR-assisted liquid-cooling for high-performance computing, developed a validated numerical model for performance and energy optimisation, and examined the potential for high-temperature heat recovery using miniature vapour-compression systems. Building on these contributions, several focused future research directions emerge.

First, future work should expand the design space of the single-phase liquid loop. The present study employed a fixed cold plate geometry and coolant configuration. However, impinging jet and/or surface feature enhancements, manifold design, coolant formulation, and pump control strategies could meaningfully influence both thermal and hydraulic behaviour. Incorporating these parameters into an extended optimisation framework would clarify the full performance envelope of hybrid liquid-VCR cooling across different processor packages.

Second, although processor package constraints were treated as fixed boundary conditions, further investigation into enhanced interface materials and surface-engineered heat spreaders would be valuable. Structured surfaces, high performance TIMs, or vapour chamber spreaders may interact differently with sub-ambient coolant temperatures. Understanding this coupling would help determine whether package-level resistances remain the dominant bottleneck even when coolant temperatures are substantially reduced.

Third, the mathematical treatment of the vapour-compression systems developed in this work provides a robust foundation for additional refinement. Opportunities include transient modelling (including duty cycling), integration of detailed compressor and heat exchanger performance maps, and consideration and optimization of more advanced evaporator and condenser heat exchanger designs.

In the context of heat recovery, the work in Chapter 5 demonstrates the feasibility of producing high-temperature heat recovery fluid using a higher pressure-ratio compressor. However, further experimental characterisation would be helpful in terms of validating the predicted performance, and practical integration with district-heating networks or DAC processes. At the facility scale, techno-economic assessment and system-level modelling would help determine the viability of deploying hybrid VCR-assisted cooling and heat recovery in future data centre environments.

A notable limitation of the present work relates to the use of R134a, a high-GWP refrigerant that is currently being phased out and is not ideal for high-temperature heat-recovery applications. Future research will investigate the

influence of alternative refrigerants - particularly low-GWP options such as R-1234yf and R-290 - on both cooling performance and energy efficiency. Given the miniature system scale and low refrigerant charge, R-290 emerges as a feasible candidate, with potential advantages in thermodynamic performance and sustainability.

Appendix A: Empirical Correlations and Heat Transfer Modelling

This appendix summarises the empirical correlations and auxiliary relations employed for evaluating air-side, liquid-side and refrigerant-side heat transfer coefficients, friction factors, and effectiveness-NTU parameters used in the heat exchanger modelling framework.

A.1. Air-Side Heat Transfer Correlation

For the louvered-fin compact heat exchangers of the condenser, the Colburn j -factor correlation proposed by Kim and Bullard [1] is used:

$$j = \text{Re}_{L_p}^{-0.487} \left(\frac{L_\alpha}{90}\right)^{0.257} \left(\frac{F_p}{L_p}\right)^{-0.13} \left(\frac{F_h}{L_p}\right)^{-0.29} \left(\frac{F_D}{L_p}\right)^{-0.235} \left(\frac{L_L}{L_p}\right)^{0.68} \left(\frac{T_p}{L_p}\right)^{-0.279} \left(\frac{F_\delta}{L_p}\right)^{-0.05} \quad (\text{A.1})$$

$$\text{Nu} = j \text{Re}_{L_p} \text{Pr}^{1/3} \quad (\text{A.2})$$

where Re_{L_p} is the Reynolds number based on louver pitch.

A.2. Single-Phase Refrigerant Flow

For both evaporator and condenser single-phase regions (superheated or subcooled), the corrected Gnielinski relation [2] is used:

$$\text{Nu} = \left\{ 1 + 7.6 \times 10^{-5} \text{Re} \left[1 - \left(\frac{d_h}{0.001167} \right)^2 \right] \right\} \times \frac{\left(\frac{f}{8} \right) (\text{Re} - 1000) \text{Pr}}{1 + 12.7 (f/8)^{0.5} (\text{Pr}^{2/3} - 1)} \quad (\text{A.3})$$

where Re is the Reynolds number, Pr is the Prandtl number, and the friction factor for smooth tubes is:

$$f = (0.79 \ln \text{Re}_d - 1.64)^{-2} \quad (\text{A.4})$$

This set of equations are valid for $0.5 \leq \text{Pr} \leq 2000$ and $2300 \leq \text{Re}_d \leq 5 \times 10^6$.

A.3. Condensation Region

Condensation heat transfer inside the condenser MPE tubes is calculated using the correlation proposed by Moser *et al.* [3]:

$$Nu = \frac{0.0994^{c_1} Re_l^{c_2} Re_{eq}^{1+0.875c_1} Pr_l^{0.815}}{(1.58 \ln Re_{eq} - 3.28)(2.58 \ln Re_{eq} + 13.7 Pr_l^{2/3} - 19.1)} \quad (A.5)$$

where,

$$c_1 = 0.126 Pr_l^{-0.448}, c_2 = -0.113 Pr_l^{-0.563} \quad (A.6)$$

and the equivalent Reynolds number is determined as,

$$Re_{eq} = \phi^{8/7} Re_l \quad (A.7)$$

$$\phi^2 = c_3 + \frac{3.24 c_4}{Fr^{0.045} We^{0.035}} \quad (A.8)$$

where the Froude number, Fr, and Weber number, We, are defined as,

$$We = \frac{G^2 d_h}{\sigma \rho_{TP}} \quad (A.9)$$

$$Fr = \frac{G^2}{g d_h \rho_{TP}^2} \quad (A.10)$$

Here, $G = \dot{m}_{ref}/A_C$ is the mass flux, g is the acceleration due to gravity, and σ is the surface tension of the refrigerant. The two-phase density, ρ_{TP} , is calculated as,

$$\rho_{TP} = \left(\frac{\chi}{\rho_g} + \frac{1-\chi}{\rho_l} \right)^{-1} \quad (A.11)$$

$$c_3 = (1-\chi)^2 + \chi^2 \left(\frac{\rho_l}{\rho_g} \right) \left(\frac{f_g}{f_l} \right) \quad (A.12)$$

$$c_4 = \chi^{0.78} (1-\chi)^{0.24} \left(\frac{\rho_l}{\rho_g} \right)^{0.91} \left(\frac{\mu_g}{\mu_l} \right)^{0.19} \left(1 - \frac{\mu_g}{\mu_l} \right)^{0.7} \quad (A.13)$$

where parameters with subscript l and g are determined by liquid and vapour phase of refrigerant respectively, and x represents the refrigerant quality at evaporator section.

To facilitate the use of the plate heat exchanger as a condenser, correlating equations for the condensation heat transfer Nusselt number provided by Yan *et al.* [4] was used;

$$Nu = 4.118 Re_{eq}^{0.4} Pr_l^{1/3} \quad (A.14)$$

with,

$$Re_{eq} = \frac{G_{eq} D_h}{\mu_l} \quad (A.15)$$

$$G_{eq} = G \left(1 - X_m + X_m \left(\frac{\rho_l}{\rho_v} \right)^{1/2} \right) \quad (A.16)$$

$$B_0 = \frac{q}{G h_{fg}} \quad (A.17)$$

here Re_{eq} is the equivalent Reynolds number, G is the mass flux, G_{eq} is the equivalent mass flux and is proposed by Akers *et al.* [5], and x is the mean vapour quality.

A.4. Evaporation Region

For refrigerant boiling within the evaporator plate heat exchanger, the correlation proposed by Ayub [6] is used to a direct expansion configuration:

$$h_{ref,sat,e} = 0.385 \left(\frac{k_l}{d_h} \right) \left[\frac{Re_l^2 La_{fg}}{L_p} \right]^{0.4124} \left(\frac{P}{P_{cr}} \right)^{0.12} \left(\frac{65}{\beta} \right)^{0.35} \quad (A.18)$$

where β is the chevron angle in the plate heat exchanger, P_e is the evaporator system pressure, P_{cr} is the critical pressure, L_p is the effective plate length of the plate heat exchanger, and $h_{ref,sat,e}$ is the latent heat of vaporisation of the refrigerant.

A.5. Liquid-Side Heat Transfer in the Primary Loop

For single-phase coolant flow in chevron-type plate heat exchangers, Okada *et al.* [7] provided the following correlation for the Nusselt number:

$$\text{Nu} = \begin{cases} 0.157 \text{Re}^{0.66} \text{Pr}^{0.4}, \beta = 60^\circ \\ 0.249 \text{Re}^{0.64} \text{Pr}^{0.4}, \beta = 45^\circ \\ 0.327 \text{Re}^{0.65} \text{Pr}^{0.4}, \beta = 30^\circ \end{cases} \quad (\text{A.19})$$

$$700 < \text{Re} < 25,000$$

For single-phase flow in a chevron-type plate heat exchanger, Martin [8] expresses the channel pressure drop as:

$$\Delta P_{\text{ch}} = 4f \frac{L_p}{D_h} \frac{G^2}{2\rho} \quad (\text{A.16})$$

where f is Fanning friction factor, which is given by:

$$f = \frac{K}{\text{Re}^m} \quad (\text{A.20})$$

where K and m are constants and are functions of the chevron angle β and are determined empirically from Heavner *et al.* [9]

A.6. Determination of Log-Mean Temperature Differences (LMTDs)

The thermal performance of the heat exchangers employed in the test loop, namely the condenser and the evaporator, was evaluated using the log-mean temperature difference (LMTD) method. This approach provides an accurate representation of the mean driving temperature difference between the hot and cold fluids when the temperature difference varies exponentially along the length of the exchanger. The LMTD for any section of a heat exchanger is expressed as,

$$\Delta T_{\text{LMTD}} = \frac{\Delta T_1 - \Delta T_2}{\ln \left(\frac{\Delta T_1}{\Delta T_2} \right)} \quad (\text{A.21})$$

where ΔT_1 = temperature difference between the hot and cold streams at one end of the section, and ΔT_2 = temperature difference between the hot and cold streams at the opposite end.

A.6.1. Condenser

The condenser was divided into three thermally distinct regions: (i) superheated vapour cooling, (ii) condensation (saturated region), and (iii) subcooled liquid cooling. The LMTD was determined individually for each region to account for the differing heat transfer mechanisms.

- Superheated Region: The refrigerant enters as a superheated vapour and is cooled to its saturation temperature. The LMTD was determined using the refrigerant and cooling air temperatures at the inlet and outlet of this segment:

$$\Delta T_1 = T_2 - T_{13}; \Delta T_2 = T_3 - T_{12} \quad (\text{A.22})$$

- Saturated (Condensation) Region: The refrigerant temperature remains constant at the saturation temperature while the cooling air temperature rises. The terminal temperature differences were calculated as:

$$\Delta T_1 = T_3 - T_{13}; \Delta T_2 = T_3 - T_{12} \quad (\text{A.23})$$

- Sub-cooled Region: The refrigerant exits the condensation zone as a saturated liquid and is further cooled to achieve the desired degree of subcooling. The LMTD for this region was similarly determined using the refrigerant and air temperatures at both ends.

$$\Delta T_1 = T_4 - T_{13}; \Delta T_2 = T_5 - T_{12} \quad (\text{A.24})$$

A.6.2. Evaporator

The evaporator was divided into two thermally distinct regions: (i) boiling (saturated region) where a liquid–vapour mixture absorbs heat at essentially constant temperature, and (ii) superheated vapour where the vapour temperature

risers above saturation. The LMTD was determined separately for each region to reflect the different heat-transfer mechanisms.

- **Boiling (Saturated) Region:** The refrigerant enters as a liquid-vapour mixture and boils at the evaporating saturation temperature T_6 while the secondary liquid is cooled from T'_8 to T_9 . With the refrigerant temperature essentially constant:

$$\Delta T_1 = T'_8 - T_6; \Delta T_2 = T_9 - T_6 \quad (\text{A.25})$$

- **Superheated Vapour Region:** After complete evaporation, the refrigerant vapour is heated from T_7 to the outlet superheat temperature T_1 , while the secondary fluid is further cooled from T_8 to T_9 . The LMTD was determined using the heated refrigerant and cooled liquid temperatures at the inlet and outlet of this segment.

$$\Delta T_1 = T_9 - T_1; \Delta T_2 = T_8 - T_7 \quad (\text{A.26})$$

A.7. Effectiveness-NTU Relations

For both evaporator and condenser regions, the heat transfer effectiveness was evaluated using the appropriate NTU-effectiveness relations. The heat capacity ratio is defined as,

$$r = \frac{C_{\min}}{C_{\max}} \quad (\text{A.27})$$

where $C_{\min} = \min(\dot{m}C_p)$ of the two fluids, and $NTU = UA/C_{\min}$. For single-phase cross-flow (both fluids unmixed) it takes the form [10],

$$\varepsilon_{sh} = 1 - e^{\left\{ \frac{NTU^{0.22}}{r} [e^{-r NTU^{0.78}} - 1] \right\}} \quad (\text{A.28})$$

For regions where phase change occurs (condensation or evaporation), the condensing/boiling fluid remains at constant temperature, and the heat capacity ratio tends to zero ($r \rightarrow 0$). In this limiting case, the effectiveness relation simplifies to [11],

$$\varepsilon_{\text{sat}} = 1 - e^{\text{NTU}} \quad (\text{A.29})$$

A.8. References


- [1] M.-H. Kim, C.W. Bullard, Air-side thermal hydraulic performance of multi-louvered fin aluminum heat exchangers, *International Journal of Refrigeration* 25 (2002) 390–400.
- [2] T.M. Adams, M.F. Dowling, S.I. Abdel-Khalik, S.M. Jeter, Applicability of traditional turbulent single-phase forced convection correlations to non-circular microchannels, *International Journal of Heat and Mass Transfer* 42 (1999) 4411–4415.
- [3] K.W. Moser, R.L. Webb, B. Na, A New Equivalent Reynolds Number Model for Condensation in Smooth Tubes, *Journal of Heat Transfer* 120 (1998) 410–417.
- [4] T.M. Adams, M.F. Dowling, S.I. Abdel-Khalik, S.M. Jeter, Applicability of traditional turbulent single-phase forced convection correlations to non-circular microchannels, *International Journal of Heat and Mass Transfer* 42 (1999) 4411–4415.
- [5] W.W. Akers, H.A. Deans, O.K. Crosser, CONDENSING HEAT TRANSFER WITHIN HORIZONTAL TUBES, *Chem. Eng. Progr. Vol: 55, Symposium Ser. No. 29* (1958).
- [6] Z.H. Ayub, Plate Heat Exchanger Literature Survey and New Heat Transfer and Pressure Drop Correlations for Refrigerant Evaporators, *Heat Transfer Engineering* (2010). <https://doi.org/10.1080/01457630304056>.
- [7] K. Okada, M. Ono, T. Tomimura, T. Okuma, H. Konno, S. Ohtani, Design and Heat Transfer Characteristics of New Plate Heat Exchanger, *Heat Transfer - Japanese Research* 1 (1972): 90-95.
- [8] H. Martin, A theoretical approach to predict the performance of chevron-type plate heat exchangers, *Chemical Engineering and Processing: Process Intensification* 35 (1996) 301–310.







- [9] Heavner, R.L., Kumar, H. and Wanniarachchi, A.S., Performance of an Industrial Plate Heat Exchanger: Effect of Chevron Angle. AIChE Symposium Series, (1993).
- [10] Y.A. Çengel, Heat and Mass Transfer: A Practical Approach, McGraw-Hill, 2007.
- [11] T.L. Bergman, F.P. Incropera, A.S. Lavine, D.P. DeWitt, Introduction to Heat Transfer, John Wiley & Sons, 2011.

Appendix B: Calibrations

B.1. Thermocouple Calibration

The following figure presents the calibration curves illustrating the predicted versus measured data recorded across all thermocouples used in this research compared to the reference probe with the associated curve equation. All probes were calibrated against a ASL F100 Precision Handheld RTD thermometer.








Figure B.1. Thermocouple calibration plots for: refrigerant loop - (a) compressor inlet (T_1), (b) compressor outlet (T_2), (c) condenser outlet (T_5), and (d) evaporator inlet (T_6); liquid loop - (e) cold plate outlet (T_8), and (f) cold plate inlet (T_9); ambient side - (g) air inlet (T_{12}), and (h) air outlet (T_{13}); TTV - (i) bottom ($T_{TTV,1}$), (j) middle ($T_{TTV,2}$) and (k) top ($T_{TTV,3}$).

B.2. Differential Pressure Sensor Calibration

The following figure presents the calibration curves illustrating the predicted vs measured data recorded across all 26PCDFA6D Honeywell differential pressure transducer used in this research, calibrated against a Digitron 2022P differential digital pressure meter. A total of three differential pressure sensors were used in this research, each across the major pressure drops in the liquid loop.




Figure B.2. Differential pressure calibration plots for: liquid loop - (a) cold plate ($\Delta P_{\text{liq,clp}}$), (b) ACHE ($\Delta P_{\text{liq,ACHE}}$), and (c) evaporator ($\Delta P_{\text{liq,e}}$).

B.3. Liquid Flow Meter Calibration

The UF25B Cynergy Ultra-Sonic (US) flow meter has its own stated error from the manufacturer, though for an extra layer of confidence, was checked a turbine flow meter (RS PRO Flow Meter 0.6 - 10 L/min; $\pm 1\%$ FS) between flow rates of 1 - 8 L/min. The calibration plot is presented below.




Figure B.3. Liquid flow meter calibration plot for liquid loop: $Q_{liq,US}$.

B.4. Calibration Certificates

The calibration certificates corresponding to the ASL F100 Precision Handheld RTD thermometer and the Digitron 2022P differential digital pressure meter, employed in the calibration of the relevant experimental apparatus, are presented below.



Industrial Temperature Sensors Ltd.
 Unit W4J, Ladystown, Toughers Business Park, Naas, Co. Kildare
 Tel: 045-408164 Fax: 045-408165 Web: www.itsirl.com email: info@itsirl.com


CERTIFICATE OF CALIBRATION

Certificate No: C226857A1	1 OF 2
Customer: Trinity College Dublin	
Instrument: Temperature Meter	
Model: ASL F100	S/N: 016219/92
Range: 0°C + 200°C	
Calibration Method: Comparison	Temperature Scale: ITS 90

Test Equipment:

- (1) DDM 900 Precision Thermometer S/N. 005100200191 Cert No. 2662
- (1a) Reference PRT Probe S/N. 09AS118 Cert No. 2662
- (2) Fluke 9172 Metrology Well S/N. B12342
- (3) Fluke 7380 High Precision Bath S/N. B14489
- (4) Lauda UB20 High Precision Bath S/N. B02001

	Reference Temperature °C	Measured Reading °C	Error °C	Measurement Uncertainty °K
CH A	-0.0014	0.1220	0.1234	0.020
CH A	20.0026	20.0060	0.0034	0.020
CH A	39.9831	40.0100	0.0269	0.020
CH A	59.9780	60.0270	0.0490	0.020
CH A	79.9826	80.0370	0.0544	0.020
CH A	99.9717	100.0240	0.0523	0.020
CH A	119.9817	120.0260	0.0443	0.020
CH A	139.9935	140.0270	0.0335	0.020
CH A	160.9776	159.9910	-0.9866	0.020
CH A	179.9628	179.9500	-0.0128	0.020
CH A	199.9648	199.9260	-0.0388	0.020

TESTED BY: Stephen Ralph


DATE: 30/09/2020

The reported expanded uncertainty is based on a standard uncertainty multiplied by a coverage factor k=2, providing a level of confidence of approximately 95%. The uncertainty evaluation has been carried out in accordance with national standards requirements. This certificate is issued in accordance with the national standards & provides traceability of measurement to recognised national standards & to units of measurement realised at the national physical laboratory or other recognised standards laboratories.



Industrial Temperature Sensors Ltd.

Unit W4J, Ladystown, Toughers Business Park, Naas, Co. Kildare
Tel: 045-408164 Fax: 045-408165 Web: www.itsirl.com email: info@itsirl.com

CERTIFICATE OF CALIBRATION

Certificate No: C126857A1

2 OF 2

Customer: Trinity College Dublin

Instrument: Temperature Meter

Model: ASL F100

S/N: 016219/92

Range: 0°C + 200°C

Calibration Method: Simulation

Temperature Scale: ITS 90

Test Equipment:

(1) DDM 900 Precision Thermometer S/N. 005100200191 Cert No. 2662

(1a) Reference PRT Probe S/N.09AS118 Cert No. 2662

(2) Fluke 9172 Metrology Well S/N. B12342

(3) Fluke 7380 High Precision Bath S/N. B14489

(4) Lauda UB20 High Precision Bath S/N. B02001

	Reference Temperature °C	Measured Reading °C	Error °C	Measurement Uncertainty °K
CH B	0.0000	0.0430	0.0430	0.020
CH B	20.0000	19.9270	-0.0730	0.020
CH B	40.0000	39.9480	-0.0520	0.020
CH B	60.0000	59.9480	-0.0520	0.020
CH B	80.0000	79.8950	-0.1050	0.020
CH B	100.0000	99.8750	-0.1250	0.020
CH B	120.0000	119.9100	-0.0900	0.020
CH B	140.0000	139.8700	-0.1300	0.020
CH B	160.0000	159.9150	-0.0850	0.020
CH B	180.0000	179.8860	-0.1140	0.020
CH B	200.0000	199.7450	-0.2550	0.020

TESTED BY: Stephen Ralph

DATE: 30/09/2020

Stephen Ralph

 Calibration Certificate - 2000 Pressure		Form 05 Rev. 2 04/04/2017 Page 1 of 1	
Registered Address: Rototherm Group Kenfig Industrial Estate Margam, Port Talbot, SA13 2PW, United Kingdom Tel: +44 (0) 1656 740 551 Registered in Wales No. 2570730 W: www.digitron.com	Customer Address: TRINITY COLLEGE DUBLIN SARA DOHERTY - DEPT MECH ENG. PARSONS BUILDING TRINITY COLLEGE DUBLIN IRELAND		
Date of Issue	22 April 2021	Certificate Number	41349
Customer Details			
Company:	TRINITY COLLEGE DUBLIN	Customer/ Works Order No.:	E12949
Customer PO Number:	PC140421GByr		
Instrument Details			
Product Code:	2082P7	Description:	2000 Series Differential Manometer
Serial Number:	4625110532	Range:	0 - 2 Bar
Customer ID No. (if req'd):	N/A	System Accuracy:	20°C to 30°C: 0.2% rdg + 0.2% fs + 1 digit; -10°C to 50°C: 0.3% rdg + 0.3% fs + 1 digit
Condition of UUT:	Good		
Traceability			
Measurements were taken in these conditions:			
Temperature (°C)	21.5		
Atmospheric Pressure (Mbar)	1012		
Measurements were taken using listed test equipment in table below:			
Test Equipment Description	Test Equipment ID	UKAS Certificate No.	
MENSOR CPC 6000	EQ.3299	20461 & N28162 to N28169	
Test Results			
Ref mbar/bar	UUT Rdg mbar/bar	PASS/FAIL	Allowable Tolerance mbar/bar
0.0	0.0	PASS	2.1
30.0	29.8	PASS	2.1
190.0	189.7	PASS	2.3
1000	999	PASS	4
1950	1949	PASS	5
The calibration was conducted in a temperature controlled environment maintained at 21 ± 2°C in accordance with the applicable calibration procedures of Rototherm Group using UKAS certified test equipment			
Tested By:	H SPAIN	Calibration Date	22 April 2021
Signature:	<i>H SPAIN</i>	Position Held	Calibration Technician
Checked by Competent Person Supervisor, Line Manger, QA, etc.		M Jones	

Template No TP 005

Template Rev 02
Company Proprietary Information

Template Rev Date: 5/04/2017

Figure B.4. The calibration certificates corresponding to the ASL F100 Precision Handheld RTD thermometer and the Digitron 2022P differential digital pressure meter

Appendix C: Properties of the Freecor - Water Coolant Mixture

C.1. Boiling Point (°C)

	Concentration of Freecor FTC (Vol%)								
	0	13.6	23	32.4	39.2	46.1	53.2	59.3	65.5
Temperature (°C)	100.0	102.7	104.2	108.0	109.3	108.4	110.5	112.3	117.5

C.2. Freezing Point (°C)

	Concentration of Freecor FTC (Vol%)								
	0	13.6	23	32.4	39.2	46.1	53.2	59.3	65.5
Temperature (°C)	-	-16.3	-25.0	-22.7	-26.0	-37.6	-50.6	-70.8	-79.8

C.3. Density (kg/m³)

	Concentration of Freecor FTC (Vol%)							
Temperature (°C)	13.6	23	32.4	39.2	46.1	53.2	59.3	65.5
20	1.0193	1.0339	1.0485	1.0585	1.0683	1.0776	1.0852	1.0924
40	1.0112	1.0250	1.0384	1.0476	1.0566	1.0653	1.0725	1.0793
60	1.0017	1.0142	1.0268	1.0354	1.0437	1.0519	1.0588	1.0654
80	0.9757	1.0012	1.0133	1.0216	1.0299	1.0378	1.0441	1.0505

C.4. Thermal conductivity (W/mK)

Temperature (°C)	Concentration of Freecor FTC (Vol%)										
	0	10	20	30	40	50	60	70	80	90	100
-50	Freez	Freez	Freez	Freez	Freez	Freez	0.427	0.417	0.389	Freez	Freez
-40	Freez	Freez	Freez	Freez	Freez	Freez	0.421	0.409	0.382	Freez	Freez
-30	Freez	Freez	Freez	Freez	Freez	0.428	0.416	0.402	0.375	0.341	Freez
-20	Freez	Freez	Freez	Freez	0.440	0.426	0.411	0.395	0.367	0.335	Freez
-10	Freez	Freez	Freez	0.462	0.441	0.423	0.406	0.388	0.360	0.329	0.317
0	Freez	0.535	0.498	0.467	0.442	0.421	0.401	0.381	0.353	0.323	0.312
10	0.575	0.546	0.506	0.471	0.443	0.418	0.396	0.374	0.346	0.317	0.306
20	0.595	0.557	0.514	0.475	0.443	0.416	0.391	0.368	0.339	0.311	0.299
30	0.613	0.567	0.521	0.479	0.443	0.413	0.386	0.361	0.331	0.305	0.292
40	0.629	0.577	0.528	0.483	0.443	0.410	0.381	0.355	0.324	0.298	0.285
50	0.642	0.586	0.535	0.487	0.443	0.408	0.376	0.349	0.317	0.292	0.277
60	0.653	0.594	0.541	0.491	0.443	0.405	0.371	0.344	0.310	0.286	0.269
70	0.661	0.602	0.547	0.494	0.442	0.403	0.366	0.338	0.303	0.280	0.261
80	0.667	0.609	0.553	0.497	0.441	0.400	0.361	0.333	0.295	0.274	0.252
90	0.671	0.616	0.558	0.500	0.441	0.398	0.356	0.328	0.288	0.268	0.243
100	Boil	0.622	0.562	0.503	0.439	0.395	0.350	0.323	0.281	0.262	0.233
110	Boil	Boil	Boil	Boil	Boil	Boil	0.345	0.318	0.274	0.256	0.223
120	Boil	Boil	Boil	Boil	Boil	Boil	Boil	Boil	0.266	0.250	0.212

C.5. Kinematic viscosity (cSt = mm²/s)

Temperature (°C)	Concentration of Freecor FTC (Vol%)								
	0	13.6	23	32.4	39.2	46.1	53.2	59.3	65.5
-30						49.14	65.78	86.22	136.6
-20					18.86	24.35	32.29	41.2	54.26
-10				8.423	10.68	13.6	17.69	22.19	28.73
0		2.871	3.897	5.313	6.63	8.361	10.58	13.21	16.59
20		1.506	1.949	2.563	3.133	3.811	4.746	5.763	7.005
40		0.943	1.186	1.485	1.77	2.13	2.59	3.051	3.656
60		0.662	0.806	0.9+79	1.145	1.343	1.593	1.867	2.181
80		0.488	0.582	0.75	0.807	0.938	1.103	1.26	1.456
100		0.389	0.451	0.536	0.62	0.709	0.805	0.916	1.04

C.6. Specific heat (kJ/kgK)

Temperature (°C)	Concentration of Freecor FTC (Vol%)										
	0	10	20	30	40	50	60	70	80	90	100
-50	Freez	Freez	Freez	Freez	Freez	Freez	2.64	2.48	2.16	Freez	Freez
-40	Freez	Freez	Freez	Freez	Freez	Freez	2.72	2.56	2.26	Freez	Freez
-30	Freez	Freez	Freez	Freez	Freez	2.95	2.81	2.64	2.35	2.31	Freez
-20	Freez	Freez	Freez	Freez	3.33	3.03	2.89	2.71	2.44	2.35	Freez
-10	Freez	Freez	Freez	3.73	3.40	3.11	2.96	2.79	2.52	2.38	2.19
0	Freez	4.05	3.91	3.77	3.46	3.18	3.03	2.85	2.60	2.42	2.24
10	4.20	4.05	3.95	3.81	3.52	3.25	3.10	2.92	2.68	2.46	2.29
20	4.19	4.06	3.98	3.85	3.57	3.31	3.16	2.98	2.75	2.50	2.34

30	4.18	4.07	4.01	3.88	3.62	3.37	3.22	3.04	2.81	2.55	2.39
40	4.18	4.08	4.03	3.91	3.66	3.42	3.27	3.09	2.87	2.60	2.44
50	4.18	4.08	4.05	3.93	3.70	3.47	3.32	3.14	2.93	2.65	2.49
60	4.18	4.09	4.06	3.95	3.73	3.52	3.37	3.19	2.99	2.70	2.54
70	4.19	4.10	4.07	3.96	3.76	3.56	3.41	3.23	3.03	2.76	2.59
80	4.20	4.11	4.08	3.97	3.78	3.59	3.44	3.27	3.08	2.82	2.64
90	4.21	4.12	4.08	3.97	3.79	3.62	3.48	3.31	3.12	2.89	2.69
100	Boil	4.13	4.07	3.97	3.80	3.65	3.50	3.34	3.15	2.95	2.74
110	Boil	Boil	Boil	Boil	Boil	Boil	3.53	3.37	3.19	3.02	2.79
120	Boil	Boil	Boil	Boil	Boil	Boil	Boil	Boil	3.21	3.09	2.84

Appendix D: Experimental Apparatus

D.1. Thermal Test Vehicle




Figure D.1. Dimensioned drawing of the Thermal Test Vehicle (TTV) (all dimensions are in mm).

D.2. Cold Plate




Figure D.2. Nexalus Enflux™ microjet impingement cold plate showing (a) assembled 3D view, (b) sectioned view, (c) jet array geometry, and (d) key geometric parameters.

D.3. Air Cooled Heat Exchanger (ACHE)



Figure D.3. Nexalus NXQ 360 fan-cooled liquid-air flat tube heat exchanger.

D.4. Fan-Fin Air-Cooler (FFAC)



Figure D.4. ARTIC Alpine Intel CPU cooler.

D.5. Aspen VCR System




Figure D.5. Aspen Liquid Chiller Module FP00039.

D.6. RIGID VCR System



Figure D.6. RIGID Plate Liquid Chiller.