

THERMODYNAMIC AND WORKLOAD
OPTIMIZATION OF DATA CENTER COOLING
INFRASTRUCTURES

THERMODYNAMIC AND WORKLOAD OPTIMIZATION OF DATA CENTER COOLING
INFRASTRUCTURES

By Rohit Gupta

A Thesis Submitted to the School of Graduate Studies in Partial Fulfillment of the Requirements
for the Degree Doctor of Philosophy

McMaster University © Copyright by Rohit Gupta, July 2021.

TITLE: Thermodynamic and workload optimization of data center cooling infrastructures

AUTHOR: Rohit Gupta,

M.Eng. (Thesis-based). (Jadavpur University, India),

B.Tech. (Maulana Abul Kalam Azad University of Technology, India).

SUPERVISOR:

Prof. Ishwar K. Puri

NUMBER OF PAGES: xxiii, 231

Abstract

The ever-growing demand for cyber-physical infrastructures has significantly affected worldwide energy consumption and environmental sustainability over the past two decades. Although the average heat load of the computing infrastructures has increased, the supportive capacity of cooling infrastructures requires further improvement. Consequently, energy-efficient cooling architectures, real-time load management, and waste heat utilization strategies have gained attention in the data center (DC) industry. In this dissertation, essential aspects of cooling system modularization, workload management, and waste-heat utilization were addressed. At first, benefits of several legacy and modular DCs were assessed from the viewpoint of the first and second laws of thermodynamics. A computational fluid dynamics simulation-informed thermodynamic energy-exergy formulation captured equipment-level inefficiencies for various cooling architectures and scenarios. Furthermore, underlying reasons and possible strategies to reduce dominant exergy loss components were suggested. Subsequently, strategies to manage cooling parameters and IT workload were developed for the DCs with rack-based and row-based cooling systems. The goal of these management schemes was to fulfill either single or multiple objectives such as energy, exergy, and computing efficiencies. Thermal models coupled to optimization problems revealed the non-trivial tradeoffs across various objective functions and operation parameters. Furthermore, the scalability of the proposed approach for a larger DC was demonstrated. Finally, a waste heat management strategy was developed for new-age infrastructures containing both air- and liquid-cooled

servers, one of the critical issues in the DC industry. Exhaust hot water from liquid-cooled servers was used to drive an adsorption chiller, which in turn produced chilled water required for the air-handler units of the air-cooled system. This strategy significantly reduced the energy consumption of existing compression chillers. Furthermore, economic and environmental assessments were performed to discuss the feasibility of this solution for the DC community. The work also investigated the potential tradeoffs between waste heat recovery and computing efficiencies.

Keywords: data center, thermal management, workload management, waste heat recovery, thermodynamic analysis, multi-objective optimization.

Acknowledgements

Firstly, I would like to thank my doctoral advisor, Prof. Ishwar K. Puri, for being a constant source of inspiration through his unparalleled scientific vision. He inspired me to take up challenges that I would never have to pursue on my own. He motivated me to be a dreamer and break free of the tethers of rationality to pursue the new. And he did all of this without ever telling me how to do it. I learned from Prof. Puri the value of being patient when the going got tough and the sense to nurture critique yet. But the most important of all is the skill to be an independent thinker, be it in an academic or a non-academic context.

Secondly, I would like to express my deepest gratitude to Prof. Douglas G. Down for his constant support and guidance in my dissertation's later phase. He always motivated me to bridge the gap between the computer science and the mechanical engineering audience to make the results more translational and easily interpretable to the target audience.

Additionally, I would like to extend my sincere appreciation to Prof. Marilyn Lightstone and Prof. James Cotton for their constant support, valuable inputs, and constructive suggestions towards improving the quality and enhancing the real-world implication of my dissertation outcome. In various stages, their challenging questions and state-of-the-art ideas have motivated me to look at the problem statement from a different perspective.

I am also deeply thankful to Dr. Souvik Pal, Dr. Suvojit Ghosh, Dr. Hosein Moazamigoodarzi, and Ms. Sahar Asgari from the Computing Infrastructure Research Center (CIRC) family where I spent the first half of my dissertation. Their support, valuable

inputs, active collaboration, and mind-teasing questions have provided me an excellent scientific orientation in the electronics cooling and data center industry.

Next, I would like to thank the Multiphysics Research Group (MuRG) family, where I spent the latter half of my Ph.D. I was lucky to perform high-quality research in biosensors and mathematical modeling of various diffusion processes as side-projects. I am thankful to Dr. Rakesh P. Sahu, Dr. Srivatsa Aithal, Ms. Sarah Mishriki, Mr. Krishna Jangid, and Mr. Tamaghna Gupta for being excellent collaborators, mentors, and labmates. I would particularly like to thank Dr. Sahu for having faith in me and encouraging me to join the SARS-CoV-2 biosensors project through the MITACS Accelerate program.

Besides, I am grateful to my prior supervisors and teachers from India, Prof. Ranjan Ganguly, Prof. Amitava Datta, and Prof. Swarnendu Sen, who motivated me to pursue doctoral research. Various formal and informal discussion with Prof. Ganguly at multiple stages of my dissertation has resulted in constructive ideas to overcome several critical hurdles of my doctoral dissertation work.

Finally, I thank my parents, who have been incredibly supportive always. Their love and constant encouragement kept me going. There had been so many tough times during this journey, but they were the ones who taught me to stay calm and patiently solve the problems that I encountered.

Rohit Gupta

Hamilton. ON

May 2021

Declaration of Academic Achievements

This dissertation was used to fulfill the requirement of a Ph.D. degree. All the research was conducted between the duration of September 2018 to August 2021. During this study, thermodynamic and computational fluid dynamics models were developed for energy-efficient data center design and operation. Furthermore, the developed energy-exergy formulations were coupled with low-complexity, zonal temperature prediction models to perform workload and cooling management in modular air-cooled DCs. These studies also captured multi-objective tradeoffs between energy, exergy, and computing efficiencies of the DCs. Finally, a waste heat management strategy by using an adsorption chiller for new-age DC rooms with both liquid- and air-cooled racks was developed. The author of the thesis is the major contributor to all the associated works presented herein.

The thesis has resulted in four manuscripts, and I am the first author on all of them.

These papers are listed below:

- **Rohit Gupta**, Sahar Asgari, Hosein Moazamigoodarzi, Souvik Pal, and Ishwar K. Puri. "Cooling architecture selection for air-cooled Data Centers by minimizing exergy destruction." *Energy* 201 (2020): 117625.
- **Rohit Gupta**, Hosein Moazamigoodarzi, SeyedMorteza MirhoseiniNejad, Douglas G. Down, and Ishwar K. Puri. "Workload management for air-cooled data centers: An energy and exergy based approach." *Energy* 209 (2020): 118485.

- **Rohit Gupta**, Sahar Asgari, Hosein Moazamigoodarzi, Douglas G. Down, and Ishwar K. Puri. "Energy, exergy and computing efficiency based data center workload and cooling management." *Applied Energy* (2021): 117050 (In press).
- **Rohit Gupta** and Ishwar K. Puri, "Waste heat recovery in a data center with an adsorption chiller: Technical and economic analysis" *Energy Conversion and Management* (2021): 114576 (In press).

Table of Contents

Abstract.....	i
Acknowledgements	iii
Declaration of Academic Achievements.....	v
Table of Contents	vii
List of Figures.....	x
List of Tables	xviii
List of Abbreviations	xxii
1 Introduction.....	1
1.1 Background.....	1
1.2 Air-cooled infrastructures.....	2
1.3 Liquid-cooled infrastructures	4
1.4 Workload management.....	6
1.5 Efficiency metrics.....	6
2 Literature Review	8
2.1 Cooling infrastructure management	8
2.2 Workload management.....	11
2.3 References	13
3 Problem statement and research objectives	23
4 Cooling architecture selection for air-cooled data centers by minimizing exergy destruction	25
4.1 Abstract.....	25
4.2 Introduction	26
4.3 Methodology.....	30

4.4	Numerical procedure	43
4.5	Results and discussion	48
4.6	Conclusion	63
4.7	Appendix: Process for exergy loss calculation.....	64
4.8	References	68
5	Energy and exergy based workload management for single-rack data centers	74
5.1	Abstract.....	74
5.2	Introduction	75
5.3	Methodology.....	78
5.4	Results and discussion	90
5.5	Conclusion	108
5.6	References	110
6	Energy, exergy, and computing efficiency based workload and cooling management for multi-rack data centers	118
6.1	Abstract.....	118
6.2	Introduction	119
6.3	Methodology.....	128
6.4	Results and discussion	144
6.5	Conclusion	166
6.6	Appendix: Pressure-flow calculation strategy using the flow network model 168	
6.7	References	171
7	Internal waste heat utilization from heterogeneous data centers using adsorption chillers	178

7.1	Abstract.....	178
7.2	Introduction	179
7.3	Description of the data center cooling infrastructure with waste heat recovery 182	
7.4	Mathematical models.....	185
7.5	Results and discussion	197
7.6	Conclusion	209
7.7	References	210
8	Conclusions.....	217
9	Future directions.....	219
10	Appendices.....	220
10.1	Performance data of the vapor compression chillers for various condenser air inlet temperatures.....	220
10.2	Detailed modeling strategy of a single server zone for the zonal temperature prediction approach.....	223
10.3	References	229

List of Figures

Figure 1.1: Generalized heat flow diagram and associated temperatures at different stages of a DC showing heat dissipation from computing chip to the facility level through a series of resistances [8]. 2

Figure 1.2: Different types of air distribution strategies and cooling geometries popularly used in air-cooled DCs. (a) perforated tiled-based uncontained system, (b) in-row cooling, (c) rack-mounted cooling, (d) overhead cooling unit based uncontained system, and (e) rear-door cooling [6, 7, 9]. 3

Figure 1.3: System-level schematic of a liquid-cooled cooled DC. (a) and (b) arrangement of components and cooling pipes inside the server, (c) supportive infrastructure by which the server-generated heat is rejected to the environment using water cooling [10]. 5

Figure 3.1: Connection and information flow between different chapters in this dissertation. 24

Figure 4.1: Thermal interactions during heat removal from the rack to the ambient. 31

Figure 4.2: Schematic representation of the geometries used to simulate four cooling architectures: (a)-(b) UFAD, (c)-(d) OHAD, (e)-(f) IRC, (g)-(h) RMCU. Blue and red dotted lines depict cold and hot airflows through the racks, respectively. 35

Figure 4.3: (a) Schematic representation of the experimental DC that has separated hot and cold chambers with 5 IT racks and 2 IRC units contained between them. In the aisle chamber, cold air exits the IRC and is drawn through the servers. Hot air exits the servers in the back aisle and is drawn into the IRC. Leakage airflow occurs through the brushes (or separators) across the aisles. Red dots indicate positions identifying the locations of

temperature probes that are placed along the halfway plane of the width of the cold ($Z = 0.06$ m) and hot ($Z = 1.22$ m) aisles. (b) top cross-sectional view showing salient airflows. The direction of leakage airflow depends on the pressure difference across the front and back chambers..... 45

Figure 4.4: Distribution of Δ in the cold and hot aisles for two different validation cases presented in Table 4.8. 48

Figure 4.5: (a) airspace exergy loss, and (b) $T_{ha} - T_{ca}$ as a function of cooling unit airflow for different geometries..... 50

Figure 4.6: Heat exchanger exergy loss as a function of cooling unit airflow for different geometries. 52

Figure 4.7: Airside and waterside components of the heat exchanger exergy loss for different geometries: (a) OHAD, (b) IRC, (c) RMCU..... 53

Figure 4.8: Influence of cooling unit airflow on the exergetic efficiencies of the heat exchangers used in the different architectures. 54

Figure 4.9: Influence of cooling unit airflow on the chiller exergy loss for different geometries. 55

Figure 4.10: Influence of cooling unit airflow on the chiller exergy loss for different geometries. 57

Figure 4.11: Contributions of different exergy loss components for the different cooling architectures with three different CRAH flowrates: (a) 3.04, (b) 4.57, and (c) 5.57 $\text{m}^3 \text{s}^{-1}$ 59

Figure 4.12: Schematic representation of exergy interactions in an air-cooled DC. 60

Figure 4.13: Dimensionless exergy loss η as a function of the dimensionless product ($Pe \times \zeta$). 62

Figure 4.14: Schematic representation of the numerical procedure for computing exergy loss in the airspace, heat exchanger and chillers..... 67

Figure 5.1: Schematic of the DC rack considered for the case study. The figure on the left depicts a three-dimensional view of the rack, while that on the right describes the salient airflows inside the enclosure. The IT rack consists of 30 1U (1U = 4.4 cm) servers and an RMCU..... 79

Figure 5.2: The process flow across different components for the DC case study. The IT rack shown in Figure 5.1 is situated within the server room interior block..... 80

Figure 5.3: A flow resistance network representation for a server enclosure equipped with thirty servers and an RMCU. For simplification, the dotted red lines represent identical sequences of servers and mechanical resistances inside the enclosure..... 81

Figure 5.4: Nonuniformity in the cold aisle temperature due to changes in T_{cw} and \dot{Q}_a for a fixed workload distribution with all servers at 50% utilization for the cases listed in Table 5.3..... 93

Figure 5.5: Non-uniformity in the cold aisle temperature due to simultaneous change in cooling parameter and workload distribution. The figure is to be interpreted in accordance with Table 5.4. 94

Figure 5.6: Effect of cooling system operating parameters on PUE and η_{2nd} for different maximum allowable temperatures in the cold aisle. The solid lines with circles represent

variations in PUE , whereas the dashed lines with square symbols represent changing η_{2nd} .
 96

Figure 5.7: Effect of cooling system operation parameters on PUE and η_{2nd} for different workload distributions. The solid lines with circles represent variations of PUE whereas dashed lines with square symbols represent changes in η_{2nd} 98

Figure 5.8: PUE minimized workload assignment and resultant cold aisle temperature distributions across the rack for different offered workloads D . The red dotted line corresponds to the maximum allowable server inlet air temperature, 27 °C. The figure should be interpreted in accordance with Table 5.5. 100

Figure 5.9: η_{2nd} maximized workload assignment and resulting cold aisle temperature distributions across the rack for different offered workloads D . The red dotted line corresponds to the maximum allowable server inlet air temperature, 27 °C. The figure should be interpreted in accordance with Table 5.6. 102

Figure 5.10: PUE and η_{2nd} as a function of LF for PUE minimization and η_{2nd} maximization problems. Solid lines with circles represent variations of PUE , whereas dashed lines with square symbols represent those for η_{2nd} . The optimization details for the figure are provided in Sections 5.3.3 and 5.3.4. 103

Figure 5.11: Pareto front of PUE and η_{2nd} obtained by solving the multi-objective optimization problem for five different LF s imposed on the DC. The symbols in the figure should be interpreted in accordance with Table 5.7 and Figure 5.12. 104

Figure 5.12: Utilization and temperature distributions for PUE and η_{2nd} aware multi-objective workload assignments. Nine salient points for the different LF marked in Figure

5.11 are considered for this diagram. The red dotted line corresponds to the maximum allowable server inlet temperature of 27 °C. The figure should be interpreted in accordance with Table 5.7. 105

Figure 5.13: Dimensionless Pareto front of PUE and η_{2nd} obtained by adopting the proposed normalization procedure. The normalization procedure nullifies the effect of LF 108

Figure 6.1: (a) Process flow layout across different components for the DC case study, and (b) representative schematic of the DC infrastructure equipped with five server racks and two IRC units. The diagram on the left shows a three-dimensional representation of the DC, whereas the one on the right shows a sectional view representing salient airflows. A server rack is divided into five zones, where each zone consists of three 2U servers (1U = 4.4 cm). There are 35 zones (25 zones in front of servers and 5 zones in front of each IRC unit) in each of the front and back chambers. 130

Figure 6.2: Generalized flow network representation (FNR) of the DC infrastructure for the case study with a total of 75 2U servers and 2 IRC units. The orange dotted lines represent identical sequences of flow resistances and servers in a simplified manner. Reproduced with permission from Ref. [22]. 132

Figure 6.3: Influence of (a) cooling parameters $T_{c,w}$ and \dot{Q}_a and (b) workload assignment on the rack height averaged cold chamber temperature distribution for the different cases listed in Table 6.4. The red dotted line in the figures denotes the maximum allowable temperature in the cold chamber. 146

Figure 6.4: Effect of load factor LF on optimum values of three objective functions: (a) Variation of COP_c and CPR_m as a function of LF , (b) Variation of COP_c and η_{ex} as a function of LF 158

Figure 6.5: Pareto optimal fronts obtained by solving the multi-objective optimization problem while satisfying $T_g = 27$ °C for different LF s. The tri-axis Pareto front is separated into three different two-dimensional Pareto fronts: (1) η_{ex} vs. COP_c , (2) CPR_m vs. COP_c , and (3) CPR_m vs. η_{ex} for better representing the trade-offs across different objectives. The “desired region” and the “undesired region” are the regimes where both the objectives are simultaneously maximized and minimized, respectively. 160

Figure 6.6: Two-dimensional distributions of zonal utilizations $u_{i,j}^z$ for nine different points marked on the Pareto fronts shown in Figure 6.5. The color bars are scaled for the extreme values of $u_{i,j}^z$, 0 and 1, respectively. This figure is to be interpreted in accordance with Table 6.10. 162

Figure 6.7: Rack height averaged cold chamber temperature distributions for the nine salient points marked on the Pareto optimal fronts presented in Figure 6.5. The “redLineTemp” corresponds to the 27 °C maximum allowable cold air temperature for servers. This figure is to be interpreted in accordance with Table 6.10 and Figure 6.6. 163

Figure 6.8: Illustration of the pod management strategy for the 100-rack DC considered. The components inside all the pods are identical and are detailed in Figure 6.1a. All pods are connected to a central load management system that presets the LF for each pod... 164

Figure 6.9: Schematic of neighboring flow resistances and pressures for a single server zone. The flow across a resistance depends on the magnitude and sign of the pressure

difference. A combination of several such server zones and the IRC units results in the generalized flow network representation shown in Figure 6.2. 169

Figure 7.1: System layout for a DC with both liquid-cooled and air-cooled racks. (a) The original system (OS), where the liquid-cooled and air-cooled systems are decoupled, and (b) the new system (NS), where these two systems are coupled through an adsorption chiller. 184

Figure 7.2: (a) Input and output parameters of the FNM used for simulating the air-cooled infrastructure with five server racks and two in-row cooling units (IRCUs). (b) Rack geometry showing the arrangement of servers and IRCUs [6]. 187

Figure 7.3: Energy savings efficiency (a, b) and annual costs (c, d) as a function of inlet water temperature of liquid-cooled racks $T_{w,i}^{lc}$ for $Q_{lc} = 300$ kW (left column) and $Q_{lc} = 150$ kW (right column). Each of the figures contains three different load sharing factors $\alpha = 0.2, 0.5,$ and 0.8 . Corresponding system parameters for the simulations are listed in Table 7. 199

Figure 7.4: Energy savings efficiency (a, b) and annual costs (c, d) as a function of the dimensionless product $\theta \times \beta$ for $Q_{lc} = 300$ kW (left column) and $Q_{lc} = 150$ kW (right column). Each of the figures contains three different load sharing factors $\alpha = 0.2, 0.5,$ and 0.8 . Corresponding system parameters for the simulations are listed in Table 8. 201

Figure 7.5: Energy savings efficiency (a, b) and annual costs (c, d) as a function of re-cooling water inlet temperature of adsorption chiller condenser $T_{c,i}^{ads}$ for (a) $Q_{lc} = 300$ kW (left column) and $Q_{lc} = 150$ kW (right column). Each of the figures contains three

different load sharing factors $\alpha = 0.2, 0.5, \text{ and } 0.8$. Corresponding system parameters for the simulations are listed in Table 9. 203

Figure 7.6: Tradeoffs between computing performance per watt (GFLOPS/W) and energy savings efficiency (ESE) as a function of inlet water temperature of liquid-cooled racks $T_{w,i}^{lc}$ for InvenSor (a) LTC30 e-plus and (b) LTC90 e-plus adsorption chiller-based waste heat recovery. 208

Figure 10.1: Chiller (a) coefficient of performance and (b) refrigeration capacity as a function of chilled water temperature for different condenser air inlet temperatures. The data is obtained from the datasheet of TRANE CGAF C25 (60 Hz) chiller with 25 Ton designed refrigeration capacity. 221

Figure 10.2: (a) Chiller coefficient of performance and (b) refrigeration capacity as a function of chilled water temperature for different condenser air inlet temperatures. The performance data is obtained from the datasheet of TRANE CGA 120 (60 Hz) chiller with 10 Ton designed refrigeration capacity. 222

Figure 10.3: Representative schematic showing salient temperatures, airflows, and energy interactions for a single server zone interacting with front (cold) and back (hot) chambers. The total volume of the server zone is identical to that of a 1U server chassis. 223

Figure 10.4: Experimentally measured (a) power consumption, (b) thermal resistance, (c) CPU temperature, and (d) χ as functions of a percentage of utilization supplied with different inlet air temperatures for HP ProLiant DL360 G5 server. 228

Figure 10.5: Total airflow of server fans as a function of inlet air temperature for HP ProLiant DL360 G5 server. 229

List of Tables

Table 4.1: Details of different cooling geometries and parameters	33
Table 4.2: Geometric considerations for CFD simulation	37
Table 4.3: Server and IT rack parameter specifications.....	38
Table 4.4: Heat exchanger characteristics [19, 40, 41].....	41
Table 4.5: Chiller parameter and specifications.....	43
Table 4.6: Mesh independence results.	44
Table 4.7: Rack operating conditions.	46
Table 4.8: Cooling unit operating conditions.....	46
Table 4.9: Decrease in exergy loss for different architectures with respect to UFAD cooling.....	59
Table 5.1: Expressions for φ_1 through φ_5 in Eq. (5.11).....	84
Table 5.2: Mathematical structure of three different optimization problems	89
Table 5.3: Case study parameters for the results presented in Figure 5.4, where variables $T_{c,w}$ and \dot{Q}_a are changed to constrain the maximum temperature in the cold aisle below 27 °C.....	91
Table 5.4: Case study parameters for the results presented in Figure 5.5, where \dot{Q}_a and workload distributions are changed to constrain the maximum temperature in the cold aisle below 27 °C.....	93
Table 5.5: Optimal solutions for the different <i>PUE</i> minimization cases presented in Figure 5.8.....	99

Table 5.6: Optimized solutions for different η_{2nd} maximization cases presented in Figure 5.9..... 101

Table 5.7: Salient points for $LF = 0.2, 0.5$ and 0.8 obtained from Figure 5.11 106

Table 6.1: Relevant state-of-the-art works in the area of thermal-aware workload and cooling management in DCs..... 122

Table 6.2: Expressions for the zonal flow interaction terms $\Omega_{1 \rightarrow 6}^f$ and $\Omega_{1 \rightarrow 6}^b$ used in Eqs. (6.9) and (6.10), respectively. 135

Table 6.3: The four different optimization problems considered. 143

Table 6.4: Parametric details of the cases presented in Figure 6.3 with $LF = 0.5$ and $\max(T_{i,j}^f) \leq 27$ °C. Red and dark green colors are mapped for two extreme values of $u_{i,j}^z$, 1 and 0, respectively. 145

Table 6.5: Influence of zonal workload assignments and cooling cycle decision variables on the three objectives for $LF = 0.5$ while satisfying the constraint $\max(T_{i,j}^f) \leq 27$ °C. For all the cases $\dot{Q}_w = 15$ GPM. Red and dark green colors are mapped for the extreme values of $u_{i,j}^z$, 1 and 0, respectively..... 148

Table 6.6: Influence of LF on decision variables and objectives while satisfying the constraint $\max(T_{i,j}^f) \leq 27$ °C. For all cases $\dot{Q}_w = 15$ GPM. 150

Table 6.7: Optimal decision variables obtained from the COP_c maximization problem for different LF . Red and dark green colors are mapped for the extreme values of $u_{i,j}^z$, 1 and 0, respectively. 152

Table 6.8: Optimal decision variables obtained from the η_{ex} maximization problem for different LF . Red and dark green colors are mapped for the extreme values of $u_{i,j}^z$, 1 and 0, respectively. 154

Table 6.9: Optimal decision variables obtained from the CPR_m maximization problem for different LF . Red and dark green colors are mapped for the extreme values of $u_{i,j}^z$, 1 and 0, respectively. 156

Table 6.10: Decision variables and objective function values for nine salient points marked on the Pareto optimal fronts presented in Figure 6.5. 161

Table 6.11: Results from the heterogeneous vs. homogeneous LF assignment across DC pods to improve CPR_m 165

Table 7.1: Performance and variable cost data for the building cooling tower used to extract heat generated by the liquid-cooled DC [25, 34]. 186

Table 7.2: Specifications and regression constants for VCR chillers with two different cooling capacities [38]. 189

Table 7.3: Specifications and regression constants for adsorption chiller with two different cooling capacities [39, 40]. 191

Table 7.4: Parameters for the cooling tower and supporting systems used to extract heat from the adsorption chiller condenser [41-43]. 193

Table 7.5: Capital expenditure (in CA\$) for the newly added equipment for waste heat recovery [41, 44]. 195

Table 7.6: Variable electricity pricing in Ontario for on-peak, mid-peak, and off-peak hours [45]. 195

Table 7.7: Parameters for the sensitivity analysis results, shown in Figure 7.3. 198

Table 7.8: Parameters for the sensitivity analysis results, shown in Figure 7.4. 200

Table 7.9: Parameters for the sensitivity analysis results, shown in Figure 7.5. 202

Table 7.10: CPP and Δm_{CO_2} for a wide range of scenarios with InvenSor LTC30 e-plus (29.5 kW) adsorption chiller retrofit. The constant parameters for these scenarios are $\dot{V}_{lc} = 100$ GPM, $Q_{lc} = 300$ kW, $Q_{ac} = 47$ kW, and $T_{amb} = 24$ °C. 205

Table 7.11: ESE, CPP, and Δm_{CO_2} for a wide range of scenarios with InvenSor LTC90 e-plus (88.5 kW) adsorption chiller retrofit. The constant parameters for these scenarios are $\dot{V}_{lc} = 100$ GPM, $Q_{lc} = 300$ kW, $Q_{ac} = 94$ kW, and $T_{amb} = 24$ °C. 207

Table 10.1: Additional specifications for TRANE CGAF C25 (60 Hz) chiller. 221

Table 10.2: Additional specification for TRANE CGA 120 (60 Hz) chiller. 222

Table 10.3: Benchmark experiments for the HP ProLiant DL360 G5 server to find unknown parameters in the thermal model. 225

List of Abbreviations

DC	Data center
UPS	Uninterrupted power supply
CRAH	Computer room air handler
CPU	Central processing unit
PUE	Power usage effectiveness
DWPE	Data center workload power efficiency
ERE	Energy reuse effectiveness
CFD	Computational fluid dynamics
UFAD	Under-floor air delivery
OHAD	Overhead air delivery
IRCU	In-row cooling unit
RMCU	Rack-mountable cooling unit
JWCM	Joint workload and cooling management
TCO	Total cost of ownership
UDF	User defined function
COP	Coefficient of performance
RANS	Reynolds' averaged Navier Stokes
RMSE	Root mean squared error
FLOPs	Floating point operations per second
DVFS	Dynamic voltage and frequency scaling

MOGA	Multi-objective genetic algorithm
VCR	Vapor compression refrigeration
FNR	Flow network representation
ASHRAE	American society of heating and refrigeration and air-conditioning engineers
ASTD	Active server temperature distribution
LF	Load factor
HPC	High-performance computing
CPR	Computing performance ratio
HRM	Heat recirculation matrix
TDP	Thermal design power
MLP	Multi-layer perceptron
CPP	Cost payback period
ESE	Energy savings efficiency
CAPEX	Capital expenditure
OPEX	Operational expenditure
LCC	Life cycle cost
ROI	Return on investment

1 Introduction

1.1 Background

Energy consumption of the cyber-physical infrastructures which support our daily computing need have risen to 4% of the global electricity demand [1]. These cyber-physical infrastructures or data centers (DCs) comprise several mission-critical equipment such as servers, data storage devices, network switches, power distribution units, and uninterrupted power supply (UPS) units. Recently, during the COVID-19 pandemic, the average load on DCs has increased by nearly 30% due to the rapid virtualization of the industrial, communication, and education sectors, resulting in a significant increment in energy consumption [2]. High-capacity computing equipment is deployed to support the computing need, which generates a substantial amount of heat. Such high heat-dissipating scenarios often cause self-shutdown and decrease thermal reliability of the chip due to steady-state and transient temperature fluctuations [3]. Such issues can cause significant loss of sensitive end-user data, exacerbate energy consumption, degrade computing equipment lifetime, or induce cascading infrastructure failure. Therefore, the DCs are equipped with a closed-loop cooling system that dissipates heat generated by the server to the surroundings. The two most popular coolants used for server heat dissipation are air and water [4]. While there have been attempts of using phase-change refrigerants, dielectric fluids, mineral oils, etc., its applicability is sparse and often questionable [4, 5]. In the subsequent sections, an overview of the air- and liquid-cooled infrastructures are presented.

1.2 Air-cooled infrastructures

One of the most widespread DC cooling techniques utilizes air as the working fluid. 95% of the commercially available computing equipment is suitable for air-cooled infrastructures. For such systems, the heat dissipated by the electronic chip is transported to the environment through the following stages [6, 7]. Cold air drawn by the server fans picks up heat rejected by the heat sink mounted on the electronics component, becoming warmer at the server exhaust. The inlet and outlet airstreams of thousands of such servers share common airspace cooled by computer room air handler (CRAH) units. Several such CRAH units collectively reject heat to the facility level cooling system, which finally releases it to the ambient. Figure 1.1 shows a generalized heat flow path from the electronic chip to the ambient and different system resistances [8].

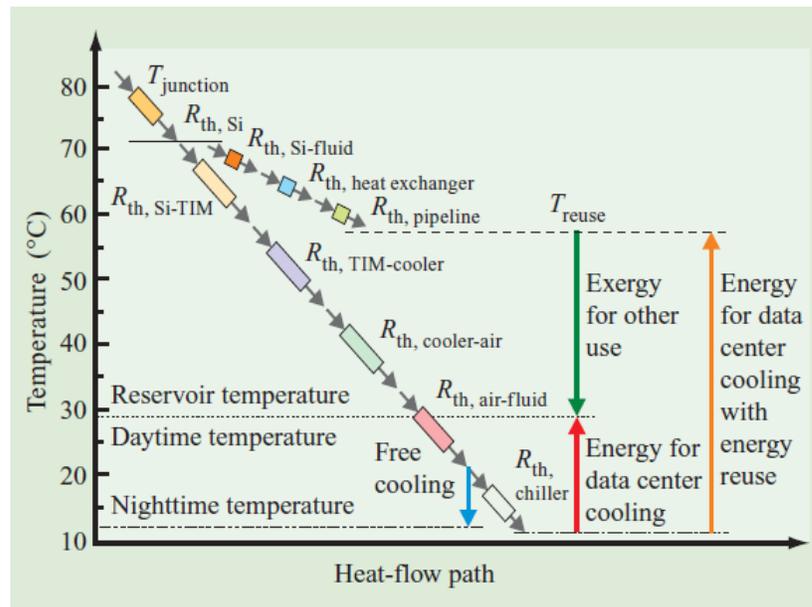


Figure 1.1: Generalized heat flow diagram and associated temperatures at different stages of a DC showing heat dissipation from computing chip to the facility level through a series of resistances [8].

An essential aspect of air-cooled DCs is the air distribution strategy. Broadly, there are three approaches reported in the literature, (1) room-based cooling, (2) row-based cooling, and (3) rack cooling [6, 7, 9]. The primary difference between strategies is the distance between the servers and the location of the CRAH unit. Further, there are different types of room cooling techniques, such as (a) perforated tile system, (b) overhead cooling units, (c) ducted return to CRAH, (d) room return to CRAH, (e) contained hot and cold aisles, (f) rear-door cooling, etc. Popular air-cooling geometries are shown in

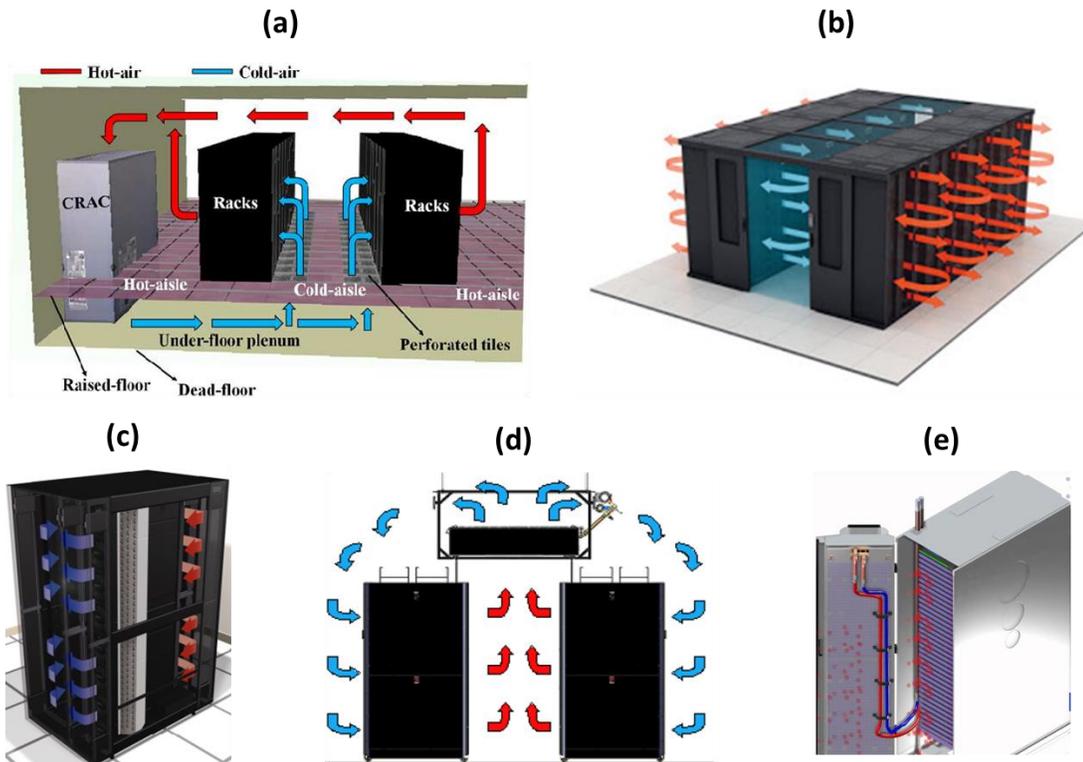


Figure 1.2: Different types of air distribution strategies and cooling geometries popularly used in air-cooled DCs. (a) perforated tiled-based uncontained system, (b) in-row cooling, (c) rack-mounted cooling, (d) overhead cooling unit based uncontained system, and (e) rear-door cooling [6, 7, 9].

1.3 Liquid-cooled infrastructures

Liquid-cooled infrastructures are adopted for DCs when the heat load of the server rack is significantly high (above 40 kW), where air-cooling cannot ensure thermally reliable operation. These systems utilize servers with liquid-cooled cold plates, microchannel heat sinks, or jet impingement heat-sinks mounted on the heat-generating electronic components. Figure 1.3 shows a representative schematic of a liquid-cooled DC infrastructure taken from [10].

In many cases, liquid-cooled DCs have proven to be much more energy-efficient than their air-cooled infrastructures [8] due to ~25% low thermal resistance of the liquid-cooled systems. Despite its potential, deploying such infrastructures require significantly high capital expenditure associated with (a) liquid-cooled server chassis, (b) rack-level coolant distribution units, (c) facility-level coolant distribution system. Moreover, the complete replacement of existing air-cooled DCs with their liquid-cooled counterpart is an impractical proposition to the DC administrator from a business perspective. In this realm, it is essential to note that several DCs administrators have partly replaced several of the existing air-cooled racks with liquid-cooled racks towards equipment upgradation. Such a transitional phase gives rise to the problem of air- and liquid- cooling hybridity inside a server room that has not been addressed to date [11]. Researchers have also leveraged the concept of phase-change cooling for DCs due to the enhanced heat transfer coefficient achieved by boiling several lab-engineered coolants [5].

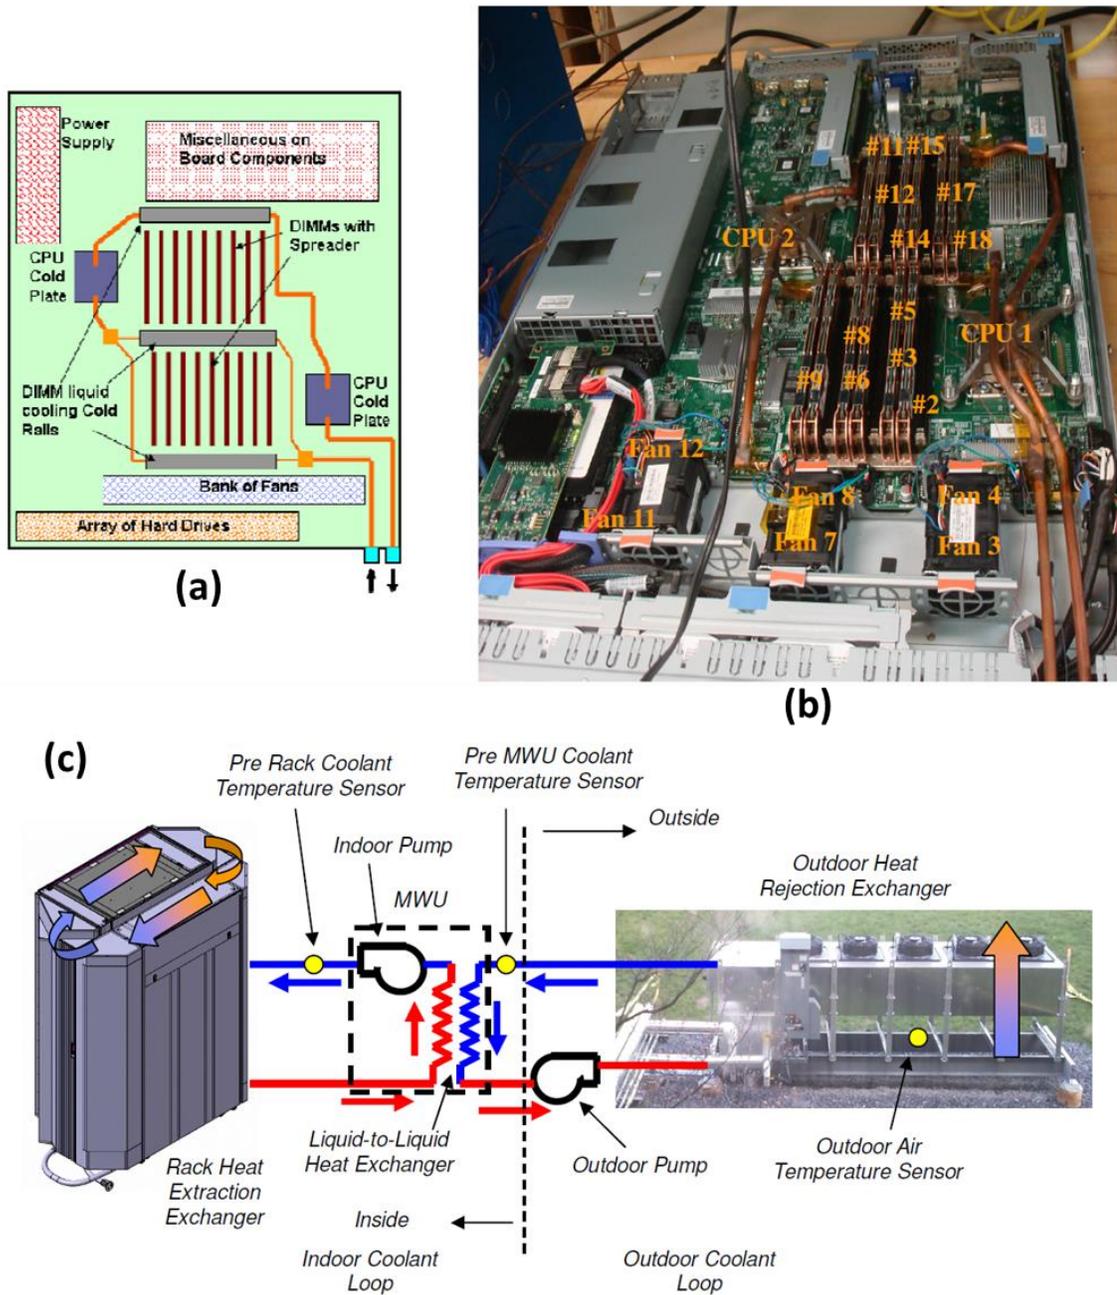


Figure 1.3: System-level schematic of a liquid-cooled cooled DC. (a) and (b) arrangement of components and cooling pipes inside the server, (c) supportive infrastructure by which the server-generated heat is rejected to the environment using water cooling [10].

1.4 Workload management

Despite having a pre-optimized cooling infrastructure at the design stage, DCs require a computing workload and cooling cycle manager during operation. The workload managers are pre-specified strategies that distribute the total incoming workload across hundreds of servers housed in the DC. Spatial heat re-distribution in DC racks can be achieved with the aid of the workload management system [12]. Besides, these are informed with additional information of the DC such as (a) temperatures at salient airspace locations, (b) coolant temperature and flowrates, (c) temperature of the central processing units (CPUs), (d) energy consumption of the cooling equipment, or (e) delay in workload processing [4] to facilitate intelligent decision making. In particular, thermally-aware workload management schemes optimizes specific goals related to thermal and energy aspects of DC infrastructures [13].

1.5 Efficiency metrics

Previously discussed cooling infrastructures, thermal management strategies, and workload managers require a common basis for comparison during both the design and operation stages. For such cases, several metrics have been reported in the DC literature are as follows [14],

The metric power usage effectiveness (*PUE*) represents the ratio of total power consumption of the DC facility to the power consumed by the computing equipment. The range of this metric is $1 \leq PUE < \infty$. An ideal but unrealistic $PUE = 1$ signifies that the DC does not consume any cooling and auxiliary power.

$$PUE = \frac{\dot{P}_{facility}}{\dot{P}_{IT}} = \frac{\dot{P}_{IT} + \dot{P}_{cooling} + \dot{P}_{auxiliary}}{\dot{P}_{IT}}. \quad (1.1)$$

The coefficient of performance in a DC is defined to characterize the energy efficiency of the cooling infrastructure use to reject the heat generated by the computing servers \dot{P}_{IT} .

$$COP = \frac{\dot{P}_{IT}}{\dot{P}_{cooling}}, \quad (1.2)$$

Data center workload power efficiency (*DWPE*) is a facility-level metric that quantifies the ratio of total computing performance supplied to the end-user to the overall power consumption of the entire DC. The computing performance of servers is often quantified by evaluating the equivalent floating-point operations per second (FLOPs). Therefore *DWPE* can also be expressed in terms of FLOPs per unit power consumption.

$$DWPE = \frac{Computing\ Performance}{\dot{P}_{cooling} + \dot{P}_{IT} + \dot{P}_{auxiliary}}. \quad (1.3)$$

Since DCs reject a significant amount of heat to the environment, which is eventually wasted, researchers have harness waste heat from DCs using different techniques. For such cases, energy reuse effectiveness (*ERE*) becomes an effective metric, a modified version of *PUE*.

$$ERE = \frac{\dot{P}_{IT} + \dot{P}_{cooling} + \dot{P}_{auxiliary} - \dot{P}_{reuse}}{\dot{P}_{IT}} = PUE - \frac{\dot{P}_{reuse}}{\dot{P}_{IT}}. \quad (1.4)$$

2 Literature Review

2.1 Cooling infrastructure management

The energy-efficient operation of DCs and airflow management is one of the critical concerns among researchers and industrial practitioners. Almost 33% of the energy consumption is attributed to the entire cooling support system, which indicates the area of potential savings [15]. The critical components of the cooling cycle, such as fans/blowers, heat exchangers, pumps, valves, chillers, and cooling towers, must be regulated in a coordinated manner to support the incoming heat load [16]. Conventionally, thermal efficiency is achieved at two stages in DCs: (a) design and (b) operation.

Researchers have investigated the influence of salient geometric parameters for perforated tile-based underfloor air delivery systems. Nada et al. [17] recommended that higher plenum heights within the range of 600 to 1080 mm and ceiling height within the range of 2000 to 3000 mm can offer more uniform airflow distribution. Additionally, cabling obstruction in the underfloor plenum can significantly reduce the air delivered to the room [6]. Later, Nada et al. [18] and Arghode et al. [19, 20] investigated the influence of perforated tile opening area, showing that perforated tile with a 25% opening area can offer a significant improvement in airflow distribution. However, using such low porosity perforated tile has its drawback. The pressure resistance will be significantly high for such cases and will prevent the CRAH supplied air from reaching the server racks [7]. This drawback of low porosity perforated tile was one of the key motivations for fan-assisted active perforated tiles proposed by Athavale et al. [21]. Additional fans are installed at the

bottom of each perforated tile, which helps overcome the pressure resistance. However, this strategy significantly increases both the capital and operational expenditure of the cooling system [9].

Another segment of the thermal management literature focuses on investigating the influence of long-distance and short-distance cooling strategies [22-24]. Long-distance cooling methods are equipped with the server racks within 10 to 20 m of the CRAH units, whereas short-distance cooling places the CRAH units with a 1 to 2 m distance of the server racks. Perforated tile-based configuration is an example of long-distance cooling [19], while row-based [23, 25] and rack-based [24, 26] cooling systems are short-distance strategies. The benefits of different cooling architecture from an energy-efficiency standpoint have been investigated by Moazamigoarzi et al. [22], which showed the benefits of row- and rack-based cooling over perforated tile systems.

Prior works show that installing aisle containment can offer significant airflow and energy efficiencies [27-29]. Aisle containments are partitions that separate the cold and hot air streams in the DC airspace. There are two possible ways in which the hot and cold air streams can mix, (1) hot air recirculation and (2) cold air bypass. The former occurs when the air supplied by CRAH airflow is significantly lower than the total suction of the server fans, while the latter occurs when the scenario is reversed. Containments can be added to different air-cooled DCs to decrease the probability of these two airspace inefficiencies. Moazamigoarzi et al. [22, 30] correlated the physics of recirculation and bypass with room-, row, and rack-based cooling architecture. This showed that row- and rack-based

systems could offer up to 29% energy savings due to reduced recirculation and bypass effects.

Researchers have examined the influence of various leakage airflow in DCs. Several possible sources of leakage are (a) backflow through air-blocking brushes [26, 31, 32], (b) reverse flow through passive servers, i.e., the servers which have their internal fans in OFF state [24, 33], (c) leakage flow through gaps under the racks [22], and (d) leakage of hot air from room to the underfloor plenum due to the venturi effect [34]. Several metrics have been reported in the literature to quantify suitable air delivery schemes, optimal geometric parameters, uniformity of airflow [35, 36]. However, most of these works do not translate the temperature-based metrics in terms of energy consumption.

A pool of literature has focused on thermodynamic analyses at a system level to improve the overall performance. Such methodology can directly quantify savings obtained from both capital and operational expenditure perspectives. Iyengar and Schmidt [37] developed an analytical thermodynamic model to investigate the energy efficiency of cooling cycle components such as blowers, pumps, chiller, and cooling tower for raised-floor DC cooling architecture. The sensitivity of the approach towards different systems parameters was also identified. Khalid et al. [38] performed equipment level thermodynamic analysis for modular DCs located at different climatic conditions. The study highlights the contribution of environmental temperature and humidity conditions towards DC energy efficiency. Patterson et al. [39] presented a comprehensive study providing perspectives on DC energy efficiency as a function of different operating temperatures.

In parallel, Shah et al. [40-43] first pointed out that energy analysis based on equipment datasheets and airflow management metrics cannot solely capture system inefficiencies in a DC cooling system. This paved a path towards the exergy destruction assessments of DC cooling systems. Early works on exergy analysis of DCs have determined component level and overall exergy losses as functions of system operating parameters [40-43]. Subsequently, exergy destruction was utilized for selecting optimal geometric parameters and operating scenarios for DCs with raised floor cooling systems [16]. Influence of rack operating conditions, load density, ambient condition, and cooling system parameters were also identified based on exergy analysis [44-46].

Recently, Silva-Llanca et al. [47-49] used computational fluid dynamics (CFD) simulation-informed exergy analysis to directly quantify the loss of available cooling energy (or exergy loss) for a real DC with raised floor and overhead configurations. This study also showed granular exergy destruction in the airspace due to hot air recirculation and cold air bypass in the DC airspace. However, exergy analysis has not yet been adopted to investigate component-level inefficiencies for a range of popular cooling infrastructures such as underfloor air delivery, overhead air delivery, row-based, and rack-based cooling systems. This gap in the literature motivates the first phase of the research work.

2.2 Workload management

Once a pre-optimized cooling infrastructure and related equipment are selected, the associated capital expenditure can be estimated. However, the related operational cost is strongly influenced by choosing cooling parameters to support specific computing loads.

Previously discussed works related to the thermodynamic analysis of DC cooling system mostly assumes a total heat load on DC and provides optimal cooling parameters. However, it was shown that the strategy by which the incoming computing workload (or the heat load) is assigned to different servers within the DCs has a strong influence on the thermal performance and energy efficiency of DCs [13, 50]. Therefore jointly managing cooling and server workload is an essential aspect of air-cooled DCs.

A large segment of the workload management literature couples thermal and energy models to distribute workload across different servers. These approaches adopt several goals such as energy efficiency [51, 52], computing performance [53], latency [54], and thermal reliability of the DC [55, 56]. The most straightforward strategy to assign workload across servers is by adopting the uniform workload management scheme, which is efficient from a computing performance perspective [57]. However, adopting this approach can drastically diminish energy efficiency and thermal reliability. CPU temperature-aware and cold air temperature-aware workload management were proposed to address the thermal reliability concerns [55, 56, 58, 59]. Since heat recirculation is also a critical inefficiency to be addressed, Tang et al. [60] developed an approach to minimize the mixing of hot and cold air in air-cooled DCs by regulating server workload. A comprehensive table is presented in the Introduction section of Chapter 6 that includes state-of-the-art work in workload management literature.

Although the approaches above somewhat reduce cooling energy expenditure while ensuring thermal safety of the servers, they do not implicitly provide a global minimum of energy consumption. To address this, energy-aware workload management schemes have

been developed. These approaches jointly regulate cooling system operating parameters and the server-level utilization distributions while minimizing the total energy expenditure. Since these schemes solve optimization problems for workload management, thermal reliability functions can be added as additional constraints.

Despite the existing workload management strategies minimize energy expenditure, they do not necessarily (1) reduce the loss of available cooling energy due to thermodynamic irreversibility, (2) fulfill multiple objectives, (3) adopt accurate thermal models for workload management, (4) develop scalable strategies for larger DCs, (5) explain the alternation of thermal field caused by choice of workload scheme, (6) advocate for the benefits of modular DCs and (7) provide operational trade-offs across different objectives.

2.3 References

- [1] E. Masanet, A. Shehabi, N. Lei, S. Smith, and J. Koomey, "Recalibrating global data center energy-use estimates," *Science*, vol. 367, pp. 984-986, 2020.
- [2] W. Strielkowski, I. Firsova, I. Lukashenko, J. Raudeliūnienė, and M. Tvaronavičienė, "Effective management of energy consumption during the COVID-19 pandemic: The role of ICT solutions," *Energies*, vol. 14, p. 893, 2021.
- [3] N. El-Sayed, I. A. Stefanovici, G. Amvrosiadis, A. A. Hwang, and B. Schroeder, "Temperature management in data centers: Why some (might) like it hot," in *Proceedings of the 12th ACM SIGMETRICS/PERFORMANCE joint international*

conference on Measurement and Modeling of Computer Systems, 2012, pp. 163-174.

- [4] A. H. Khalaj and S. K. Halgamuge, "A Review on efficient thermal management of air-and liquid-cooled data centers: From chip to the cooling system," *Applied energy*, vol. 205, pp. 1165-1188, 2017.
- [5] B. B. Kanbur, C. Wu, S. Fan, and F. Duan, "System-level experimental investigations of the direct immersion cooling data center units with thermodynamic and thermoeconomic assessments," *Energy*, vol. 217, p. 119373, 2021.
- [6] W.-X. Chu and C.-C. Wang, "A review on airflow management in data centers," *Applied Energy*, vol. 240, pp. 84-119, 2019.
- [7] H. Lu, Z. Zhang, and L. Yang, "A review on airflow distribution and management in data center," *Energy and Buildings*, vol. 179, pp. 264-277, 2018.
- [8] T. Brunschwiler, B. Smith, E. Ruetsche, and B. Michel, "Toward zero-emission data centers through direct reuse of thermal energy," *IBM Journal of Research and Development*, vol. 53, pp. 11: 1-11: 13, 2009.
- [9] K. Zhang, Y. Zhang, J. Liu, and X. Niu, "Recent advancements on thermal management and evaluation for data centers," *Applied thermal engineering*, vol. 142, pp. 215-231, 2018.
- [10] P. R. Parida, T. Chainer, M. Iyengar, M. David, M. Schultz, M. Gaynes, *et al.*, "System-level design for liquid cooled chiller-less data center," in *ASME*

International Mechanical Engineering Congress and Exposition, 2012, pp. 1127-1139.

- [11] K. Ebrahimi, G. F. Jones, and 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*, vol. 31, pp. 622-638, 2014.
- [12] S. MirhoseiniNejad, G. Badawy, and D. G. Down, "Holistic thermal-aware workload management and infrastructure control for heterogeneous data centers using machine learning," *Future Generation Computer Systems*, vol. 118, pp. 208-218, 2021.
- [13] Z. Abbasi, G. Varsamopoulos, and S. K. Gupta, "Thermal aware server provisioning and workload distribution for internet data centers," in *Proceedings of the 19th ACM International Symposium on High Performance Distributed Computing*, 2010, pp. 130-141.
- [14] V. D. Reddy, B. Setz, G. S. V. Rao, G. Gangadharan, and M. Aiello, "Metrics for sustainable data centers," *IEEE Transactions on Sustainable Computing*, vol. 2, pp. 290-303, 2017.
- [15] M. Iyengar and R. Schmidt, "Analytical modeling for thermodynamic characterization of data center cooling systems," *Journal of Electronic Packaging*, vol. 131, 2009.
- [16] K. Fouladi, A. P. Wemhoff, L. Silva-Llanca, K. Abbasi, and A. Ortega, "Optimization of data center cooling efficiency using reduced order flow modeling

within a flow network modeling approach," *Applied Thermal Engineering*, vol. 124, pp. 929-939, 2017.

- [17] S. Nada and M. Said, "Comprehensive study on the effects of plenum depths on air flow and thermal managements in data centers," *International Journal of Thermal Sciences*, vol. 122, pp. 302-312, 2017.
- [18] S. Nada, K. Elfeky, A. M. Attia, and W. Alshaer, "Experimental parametric study of servers cooling management in data centers buildings," *Heat and Mass Transfer*, vol. 53, pp. 2083-2097, 2017.
- [19] V. K. Arghode and Y. Joshi, "Experimental investigation of air flow through a perforated tile in a raised floor data center," *Journal of electronic packaging*, vol. 137, 2015.
- [20] V. K. Arghode and Y. Joshi, "Modeling strategies for air flow through perforated tiles in a data center," *IEEE Transactions on Components, Packaging and Manufacturing Technology*, vol. 3, pp. 800-810, 2013.
- [21] J. Athavale, Y. Joshi, and M. Yoda, "Experimentally validated computational fluid dynamics model for data center with active tiles," *Journal of Electronic Packaging*, vol. 140, 2018.
- [22] H. Moazamigoodarzi, P. J. Tsai, S. Pal, S. Ghosh, and I. K. Puri, "Influence of cooling architecture on data center power consumption," *Energy*, vol. 183, pp. 525-535, 2019.

- [23] H. Moazamigoodarzi, R. Gupta, S. Pal, P. J. Tsai, S. Ghosh, and I. K. Puri, "Modeling temperature distribution and power consumption in IT server enclosures with row-based cooling architectures," *Applied Energy*, vol. 261, p. 114355, 2020.
- [24] H. Moazamigoodarzi, S. Pal, S. Ghosh, and I. K. Puri, "Real-time temperature predictions in it server enclosures," *International Journal of Heat and Mass Transfer*, vol. 127, pp. 890-900, 2018.
- [25] J. Priyadumkol and C. Kittichaikarn, "Application of the combined air-conditioning systems for energy conservation in data center," *Energy and buildings*, vol. 68, pp. 580-586, 2014.
- [26] H. Moazamigoodarzi, S. Pal, D. Down, M. Esmalifalak, and I. K. Puri, "Performance of a rack mountable cooling unit in an IT server enclosure," *Thermal Science and Engineering Progress*, vol. 17, p. 100395, 2020.
- [27] S. Khalili, M. I. Tradat, K. Nemati, M. Seymour, and B. Sammakia, "Impact of tile design on the thermal performance of open and enclosed aisles," *Journal of Electronic Packaging*, vol. 140, 2018.
- [28] D. W. Demetriou and H. E. Khalifa, "Optimization of enclosed aisle data centers using bypass recirculation," *Journal of Electronic Packaging*, vol. 134, 2012.
- [29] H. S. Erden, M. Koz, M. T. Yildirim, and H. E. Khalifa, "Optimization of enclosed aisle data centers with induced CRAH bypass," *IEEE Transactions on Components, Packaging and Manufacturing Technology*, vol. 7, pp. 1981-1989, 2017.
- [30] H. Moazamigoodarzi, "Distributed cooling for data centers: Benefits, performance evaluation and prediction tools," Ph.D. Thesis, McMaster University, 2019.

- [31] A. Almoli, A. Thompson, N. Kapur, J. Summers, H. Thompson, and G. Hannah, "Computational fluid dynamic investigation of liquid rack cooling in data centres," *Applied energy*, vol. 89, pp. 150-155, 2012.
- [32] S. Asgari, H. Moazamigoodarzi, P. J. Tsai, S. Pal, R. Zheng, G. Badawy, *et al.*, "Hybrid surrogate model for online temperature and pressure predictions in data centers," *Future Generation Computer Systems*, vol. 114, pp. 531-547, 2021.
- [33] S. MirhoseiniNejad, H. Moazamigoodarzi, G. Badawy, and D. G. Down, "Joint data center cooling and workload management: A thermal-aware approach," *Future Generation Computer Systems*, vol. 104, pp. 174-186, 2020.
- [34] M. Xie, J. Wang, and J. Liu, "Evaluation metrics of thermal management in data centers based on exergy analysis," *Applied Thermal Engineering*, vol. 147, pp. 1083-1095, 2019.
- [35] A. Capozzoli, G. Serale, L. Liuzzo, and M. Chinnici, "Thermal metrics for data centers: A critical review," *Energy Procedia*, vol. 62, pp. 391-400, 2014.
- [36] A. Capozzoli, M. Chinnici, M. Perino, and G. Serale, "Review on performance metrics for energy efficiency in data center: The role of thermal management," in *International Workshop on Energy Efficient Data Centers*, 2014, pp. 135-151.
- [37] R. Steinbrecher and R. Schmidt, "Data center environments," *ASHRAE J*, vol. 42, p. e49, 2011.
- [38] R. Khalid, A. P. Wemhoff, and Y. Joshi, "Energy and exergy analysis of modular data centers," *IEEE Transactions on Components, Packaging and Manufacturing Technology*, vol. 7, pp. 1440-1452, 2017.

- [39] M. K. Patterson, "The effect of data center temperature on energy efficiency," in *2008 11th Intersociety Conference on Thermal and Thermomechanical Phenomena in Electronic Systems*, 2008, pp. 1167-1174.
- [40] A. J. Shah, V. P. Carey, C. E. Bash, C. D. Patel, and R. K. Sharma, "Exergy Analysis of Data Center Thermal Management Systems," in *Energy Efficient Thermal Management of Data Centers*, Y. Joshi and P. Kumar, Eds., ed Boston, MA: Springer US, 2012, pp. 383-446.
- [41] A. J. Shah, V. P. Carey, C. E. Bash, and C. D. Patel, "Exergy-based optimization strategies for multi-component data center thermal management: Part II—application and validation," in *International Electronic Packaging Technical Conference and Exhibition*, 2005, pp. 215-224.
- [42] A. J. Shah, V. P. Carey, C. E. Bash, and C. D. Patel, "Exergy-based optimization strategies for multi-component data center thermal management: Part I—analysis," in *International Electronic Packaging Technical Conference and Exhibition*, 2005, pp. 205-213.
- [43] A. J. Shah, V. P. Carey, C. E. Bash, and C. D. Patel, "Exergy analysis of data center thermal management systems," *Journal of Heat Transfer*, vol. 130, 2008.
- [44] A. J. Díaz, R. Cáceres, J. M. Cardemil, and L. Silva-Llanca, "Energy and exergy assessment in a perimeter cooled data center: The value of second law efficiency," *Applied Thermal Engineering*, vol. 124, pp. 820-830, 2017.

- [45] X. Qian, Z. Li, and Z. Li, "Entransy and exergy analyses of airflow organization in data centers," *International Journal of Heat and Mass Transfer*, vol. 81, pp. 252-259, 2015.
- [46] A. Bhalerao, K. Fouladi, L. Silva-Llanca, and A. P. Wemhoff, "Rapid prediction of exergy destruction in data centers due to airflow mixing," *Numerical Heat Transfer, Part A: Applications*, vol. 70, pp. 48-63, 2016.
- [47] L. Silva-Llanca, A. Ortega, K. Fouladi, M. del Valle, and V. Sundaralingam, "Determining wasted energy in the airside of a perimeter-cooled data center via direct computation of the Exergy Destruction," *Applied Energy*, vol. 213, pp. 235-246, 2018.
- [48] L. Silva-Llanca, M. del Valle, and A. Ortega, "The effectiveness of data center overhead cooling in steady and transient scenarios: comparison of downward flow to a cold aisle versus upward flow from a hot aisle," in *ASME 2015 International Technical Conference and Exhibition on Packaging and Integration of Electronic and Photonic Microsystems collocated with the ASME 2015 13th International Conference on Nanochannels, Microchannels, and Minichannels*, 2015.
- [49] L. Silva-Llanca, M. Del Valle, A. Ortega, and A. J. Díaz, "Cooling effectiveness of a data center room under overhead airflow via entropy generation assessment in transient scenarios," *Entropy*, vol. 21, p. 98, 2019.
- [50] Y. Li, X. Wang, P. Luo, and Q. Pan, "Thermal-aware hybrid workload management in a green datacenter towards renewable energy utilization," *Energies*, vol. 12, p. 1494, 2019.

- [51] S. M. M. Nejad, G. Badawy, and D. G. Down, "EAWA: Energy-aware workload assignment in data centers," in *2018 International Conference on High Performance Computing & Simulation (HPCS)*, 2018, pp. 260-267.
- [52] A. A. Khan, M. Zakarya, and R. Khan, "Energy-aware dynamic resource management in elastic cloud datacenters," *Simulation modelling practice and theory*, vol. 92, pp. 82-99, 2019.
- [53] D. C. Price, M. A. Clark, B. R. Barsdell, R. Babich, and L. J. Greenhill, "Optimizing performance-per-watt on GPUs in high performance computing," *Computer Science-Research and Development*, vol. 31, pp. 185-193, 2016.
- [54] S. Ghorbani, Z. Yang, P. B. Godfrey, Y. Ganjali, and A. Firoozshahian, "Drill: Micro load balancing for low-latency data center networks," in *Proceedings of the Conference of the ACM Special Interest Group on Data Communication*, 2017, pp. 225-238.
- [55] Y. Bai and L. Gu, "Chip temperature-based workload allocation for holistic power minimization in air-cooled data center," *Energies*, vol. 10, p. 2123, 2017.
- [56] Y. Bai, L. Gu, and X. Qi, "Comparative study of energy performance between chip and inlet temperature-aware workload allocation in air-cooled data center," *Energies*, vol. 11, p. 669, 2018.
- [57] M. Harchol-Balter, *Performance modeling and design of computer systems: queueing theory in action*: Cambridge University Press, 2013.

- [58] J. D. Moore, J. S. Chase, P. Ranganathan, and R. K. Sharma, "Making Scheduling" Cool": Temperature-Aware Workload Placement in Data Centers," in *USENIX annual technical conference, General Track*, 2005, pp. 61-75.
- [59] L. Wang, S. U. Khan, and J. Dayal, "Thermal aware workload placement with task-temperature profiles in a data center," *The Journal of Supercomputing*, vol. 61, pp. 780-803, 2012.
- [60] Q. Tang, S. K. S. Gupta, and G. Varsamopoulos, "Energy-efficient thermal-aware task scheduling for homogeneous high-performance computing data centers: A cyber-physical approach," *IEEE Transactions on Parallel and Distributed Systems*, vol. 19, pp. 1458-1472, 2008.

3 Problem statement and research objectives

Out of the various cooling configurations discussed in the literature review section, the perforated tile, rear-door cooling, and overhead cooling systems are popular in room-based configurations. While the in-row cooling (IRC) and rack-mounted cooling (RMC) are becoming more popular in recently deployed high load density DCs up to 30 kW of heat load per rack. Although the literature provides several temperature-based and energy-based metrics for assessing the relative benefits of these configurations, the reason for the cooling inefficiency of these geometries is not well understood. A second-law-based assessment can offer important granular information of cooling inefficiency (or exergy destruction) in the DC airspaces and various equipment in this realm. Therefore, the first objective (Chapter 4) is to quantify exergy destruction and exergy efficiencies of different DC components such as servers, airspace, heat exchanger, and chillers. The proposed methodology combined turbulent computational fluid dynamics simulation of various DC cooling geometries and informed thermodynamic energy-exergy balance formulation.

Exergy efficient modular geometries obtained from the first objective, i.e., IRC and RMC, require further optimization during operation to fulfill multiple goals of a DC administrator. Conventionally, the industrial DCs adopt joint workload and cooling management (JWCM) strategies to meet the load demand providing a thermally reliable environment to the IT equipment. However, for IRC and RMC geometries, such JWCM approaches addressing essential objectives such as energy, exergy, and computing efficiencies have not yet been developed. This motivates the second and third objectives

(Chapters 5 and 6) in which the thermally-aware, multi-objective JWCM strategies for RMC and IRC architecture were proposed. The methodology combines reduced-order spatiotemporal temperature prediction models with formulations of energy exergy and computing efficiencies. In this objective, the essential tradeoffs across different objective functions, workload distribution, IT load factor, and thermal reliability guidelines were investigated to enhance overall DC performance.

Finally, in the fourth stage (Chapter 7), a crucial cooling heterogeneity discussed in Section 1.2 has been addressed in newly retrofitted DCs that contains a fraction of racks with IRC configuration, whereas the other racks are water-cooled. A waste heat management strategy to couple these racks using thermally driven adsorption chilled was developed in this stage. The waste heat extracted from liquid-cooled racks was used to drive the adsorption chiller, which, in turn, produced chilled water for the IRC units for the air-cooled racks. The case study adopted energy, exergy, economic, and environmental impact analysis to examine the feasibility of the proposed solution. Additionally, the work revealed non-trivial tradeoffs between waste heat recovery and computing efficiency. To aid the reader, Figure 3.1 shows a comprehensive flow of information in this thesis.

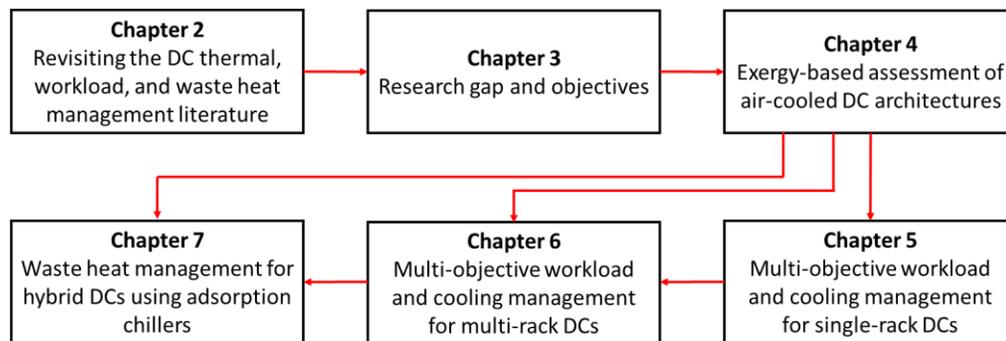


Figure 3.1: Connection and information flow between different chapters in this dissertation.

4 Cooling architecture selection for air-cooled data centers by minimizing exergy destruction

This chapter is reproduced from *Cooling architecture selection for air-cooled data centers by minimizing exergy destruction*, *Energy* 201 (2020): 117625, **Rohit Gupta**, Sahar Asgari, Hosein Moazamigoodarzi, Souvik Pal, and Ishwar K. Puri, <https://doi.org/10.1016/j.energy.2020.117625>, Copyright © 2021, with permission from Elsevier. The author of this thesis is the first author and the primary contributor to this publication.

4.1 Abstract

Air-cooled Data Centers (DCs) require effective thermal management of the servers. This can be accomplished by implementing new cooling architectures. Nearly 33% of overall energy consumption is attributed to the cooling infrastructure, which indicates the importance of the specific cooling configuration. Our objective is to compare four emerging and traditional DC cooling architectures, (a) in-row cooling, (b) rack-mountable cooling (RMC), (c) underfloor air delivery (UFAD), and (d) overhead air delivery. Since a first law-based energy analysis of a DC cooling architecture seldom considers irreversibility and component level inefficiency, an exergy-based analysis provides an alternate basis for assessment. We propose a methodology that combines Computational Fluid Dynamics simulations with thermodynamic energy and exergy balances to determine the exergy loss in different components in DC. A dimensionless parameter is identified to characterize the exergy loss as a function of the Peclet number and the dimensionless dead

state temperature ratio. The architecture containing RMC units has the lowest exergy loss. The chiller loss constitutes up to 55% of the overall exergy loss. This analysis facilitates better decision-making and design choices for air-cooled DCs on the basis of minimization of thermodynamic irreversibility with the purpose of lowering energy waste.

4.2 Introduction

Data Centers (DC) are meticulously designed facilities that house complex IT equipment (ITE), which includes servers, network switches, routers, storage systems, and network firewalls. Due to the essential services that they enable, e.g., banking, healthcare, and defense, DCs are often referred to as mission-critical facilities [1]. With hardware and software advances, the computational power of ITE is rapidly increasing, necessitating improvements in the cooling provided to it during runtime. Almost one-third of the total power consumption in a DC is associated with cooling, which provides an opportunity to identify potential areas for savings [2, 3]. Although liquid-cooled DCs are much more energy-efficient, they require a complex plumbing loop and impose reliability concerns [4, 5]. Hence, most DCs are designed to house air-breathing servers. These servers draw in cold air for cooling and exhaust the heated air. The turbulent airflow in a DC is a dynamically changing vector field. Consequently, the choice of cooling architecture has a significant influence on the system energy efficiency [6, 7].

Broadly, four cooling architectures are employed, (1) underfloor air delivery (UFAD), (2) overhead air delivery (OHAD), (3) in-row cooling (IRC), and (4) rack-mountable cooling units (RMCU). The first two systems are more common, while the latter two are relatively new. There are two primary air distribution inefficiencies in a DC, hot

air recirculation and cold air bypass [8]. A poorly managed DC airspace has a cascading detrimental effect on other cooling components, such as the computer room air handler (CRAH) unit, chillers, and cooling towers, increasing overall cooling energy consumption in a DC for a specific IT load.

The literature provides energy assessments of DC cooling architectures based on the first law of thermodynamics [6, 7] and uses temperature-based metrics, such as supply heat index *SHI*, return heat index *RHI*, rack temperature index *RTI*, rack cooling index *RCI*, negative pressure ratio *NP*, recirculation ratio *R*, bypass ratio *BP*, and balance ratio *BAL* to characterize airspace inefficiencies [9-11]. Analyses based on the first law of thermodynamics only account for the amount of energy spent and not the irreversibility due to entropy generation. Since they cannot quantify the waste of available energy, these first law methods are insufficient for developing designs for effective thermal management of air-cooled DCs. To better understand cooling inefficiencies, an analysis based on the second law of thermodynamics that characterizes exergy destruction in a DC becomes essential.

The transport of heat in an air-cooled DC occurs across multiple components, e.g., the CPU, copper heat sink on the CPU, airflow through the heat sink, air to liquid heat exchanger, and heat transfer from the liquid to refrigerant and refrigerant to ambient air. Since exergy destruction during these stages is associated with the additional power consumed for cooling, it becomes possible to identify system design features that lead to cooling inefficiencies.

The literature on exergy-based assessments that can guide improvements in DC design is sparse. It includes full system modeling of the exergy loss in a DC for a conventional UFAD configuration using a thermodynamic energy and exergy balance methodology known as the indirect method [12-17]. These investigations identify how exergy destruction occurs in each DC component due to varying rack operation conditions, IT load density, cooling unit operating parameters, and chiller and cooling tower operation. The direct method to determine exergy destruction requires fluid dynamics solutions. Here, computational fluid dynamics (CFD) based solutions of Reynolds averaged Navier Stokes (RANS) equations can be used to resolve the flow inside a DC airspace, e.g., to determine the exergy loss in the UFAD and OHAD architectures that are popular in legacy DCs [18-20]. The former type of architecture is an underfloor perforated tile-based cooling system, whereas the latter type employs an overhead cooling system. For these architectures, since the volume of air handled is prodigious and occurs over large flow path lengths, the corresponding pressure drops introduce inefficiencies in the system. Both architectures are also prone to hot air recirculation and cold air bypass.

Although legacy architectures employ room-based cooling architectures, the RMCU is promising for high-density computing infrastructure [21-23]. Such a scheme consists of a rack-mountable cooling system placed inside each rack that has separated hot and cold chambers. A version of the RMCU suitable for a high density scalable modular DC is the IRC that simultaneously delivers cooling air to several racks to reduce hot-spots [24, 25]. The IRC architecture is an enclosed row-based cooling solution that provides cold air to several IT racks stacked beside one another. Placing the cooling units nearer to the

heat sources, i.e., the servers, reduces the airflow path length, which in turn reduces the adverse effects due to pressure drops. Doing so also reduces hot and cold air mixing. An exergy-based assessment of the improvement in energy consumption by a modular (i.e., RMCU and IRC) DC as compared to a legacy DC is yet unavailable. Hence, there is a lack of guidance for DC designers for selecting a suitable cooling architecture. This leads to cooling overdesign, often by a factor greater than two, producing energy waste and considerably increasing the total cost of ownership (TCO).

We develop a methodology that compares different cooling architectures based on their contributions to exergy destruction. The airspace exergy loss parameters are obtained with a direct method, whereas an indirect method was applied to quantify exergy losses in servers, CRAH units, and the chiller. This hybrid approach provides full system exergy modeling, i.e., extending from servers to chillers while reducing computational expense. The method is used to compare four different cooling schemes (UFAD, OHAD, IRC, and RMCU). For each geometry, seven different scenarios are considered by varying the operation of the DC cooling unit, and the resulting impact on component level exergy loss is identified. Finally, a dimensionless parameter to characterize overall exergy destruction in DC is identified.

The novelty of the study lies in the development of a hybrid method that characterizes exergy destruction in a DC and thus comparing different cooling architectures. A dimensionless parameter is a facile tool that provides designers with guidance on how to minimize overall exergy loss in a DC. The method identifies (a) designs for optimal cooling, (b) component level inefficiencies, (c) favorable air delivery schemes,

(d) the impact of operating conditions on cooling energy wastage, and (e) it also minimizes exergy destruction.

Our objectives are to (1) demonstrate that the exergy destruction minimization approach is suitable for comparing the performance of DC cooling systems, (2) investigate the component level irreversibility in servers, DC airspace, CRAH unit, and chillers, (3) compare different cooling architectures for which we have selected the UFAD, OHAD, IRC, and RMCU configurations, (4) investigate the influence of cooling unit airflow on the exergy destruction that occurs in different DC components, (5) propose strategies to minimize exergy loss in each component, and (6) identify a dimensionless parameter that characterizes the overall exergy loss in a DC.

4.3 Methodology

We consider a 63 kW air-cooled DC that uses CRAH units and air-cooled chillers to reject the heat from the DC airspace. Figure 4.1 presents the thermal interactions across the DC components. Servers generate heat that is removed by cold air supply, and the warmer server exhaust is returned to the CRAH units, which employ air to water heat exchanger to remove heat from the airside and transfer it to the waterside. A pump-driven chilled water circulation loop connects the waterside of the CRAH units to an air-cooled chiller that rejects heat to the ambient. The chiller has a nominal design capacity of 88 kW [26]. While other options for the chiller are available for DCs with higher capacities, such as a water-cooled chiller that rejects heat to the cooling tower [17], our selection is nevertheless representative for comparing component level losses (e.g., in servers, airspace, CRAH,

pumps, and chillers) and the overall exergy destruction due to changes in the cooling architecture.

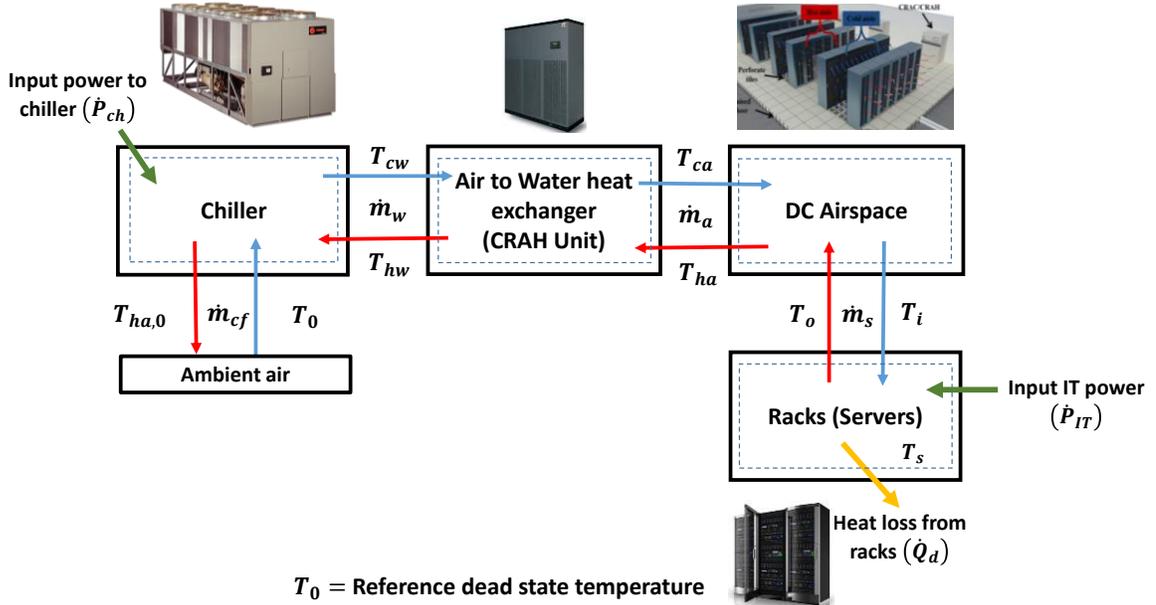


Figure 4.1: Thermal interactions during heat removal from the rack to the ambient.

In an air-cooled DC, exergy destruction depends on the (1) cooling architecture, (2) flowrate of cold air delivery, and (3) the mismatch between server suction temperature and cooling unit setpoint. By decomposing the overall exergy destruction for its principal components, we obtain the exergy destruction in the (1) servers, (2) heat exchanger inside the CRAH unit, (3) fans inside the CRAH unit, (4) DC airspace, (5) chilled water pumps, and (6) chillers. Items 3 and 5 consume 4% and 7% of the total cooling power, respectively. The efficiencies of the pumps and fans used in DC are assumed to be 75% [7, 17, 27], i.e., exergy destruction in pumps and fans is negligible (of the order of 2 to 3 %) compared to the magnitude of overall exergy loss. Therefore, we only determine the exergy loss components associated with the (a) servers, (b) airspace, (c) heat exchanger inside CRAH units, and (d) chillers for different cooling schemes.

To determine the exergy destruction for these components, the following conditions must be determined or assumed, (1) thermodynamic dead state temperature (assumed at a constant value of 295.15 K), (2) total IT load, which is 63 kW in our case, (3) airspace field variables, i.e., pressure, temperature, velocities, and turbulence parameters, (4) CRAH airflow, (5) CRAH setpoint, (6) airside return temperature to CRAH units, (7) supply chilled water temperature to CRAH, (8) return water temperature from CRAH, (9) chilled water flow rate, (10) ambient air temperature supplied to the chiller (assumed to equal to the thermodynamic dead state temperature), (11) heat rejection temperature to the ambient from chiller, and (12) airflows of chiller fans.

To compare operation over a range of operating parameters for different cooling architectures, the airflow through the CRAH unit is varied, as shown in Table 4.1. Following the appropriate ASHRAE guideline for the safe operation of servers, the maximum server intake temperatures are held below 26.5 °C [28, 29] for all the scenarios simulated. Since the cooling unit supply temperature is specified, the resulting airflow and return temperatures are determined using CFD for different air delivery schemes. The airside and waterside parameters are coupled by using the $\epsilon - NTU$ method [6, 17, 30] for heat exchangers within the CRAH units for different cooling architectures.

Table 4.1: Details of different cooling geometries and parameters

IT load (kW)	Geometry	Cold air supply temperature (°C)	Total volume flow rate of air ($m^3 s^{-1}$)	Maximum temperature in the cold aisle (°C)
63	UFAD	18	2.53	26.5
			3.04	
	OHAD		3.54	
			4.05	
	IRC		4.56	
	RMCU		5.06	
			5.57	

4.3.1 Geometries and physical considerations for CFD simulation

The turbulent flow field in the DC airspace is resolved through CFD simulations using ANSYS Fluent 18.0 software. The RANS equations are coupled with the energy equation and a standard $k - \varepsilon$ model [6, 18-20, 31-33] for steady-state, incompressible flow, and a Boussinesq approximated body force to determine the velocity, temperature, and turbulence fields. Figure 4.2 provides a schematic representation of the different air delivery schemes considered, which are (a) UFAD, (b) OHAD, (c) IRC, and (d) RMCU. For all architectures, we consider 6 IT racks of standard dimensions ($2 \times 1 \times 0.6 m^3$), each consisting of 40 1U (Model: HP ProLiant DL360 G5) servers. Each rack has an IT load of 10.5 kW distributed over 40 servers and an airflow of $0.57 m^3 s^{-1}$, where each server has an airflow of $0.01415 m^3 s^{-1}$ [21].

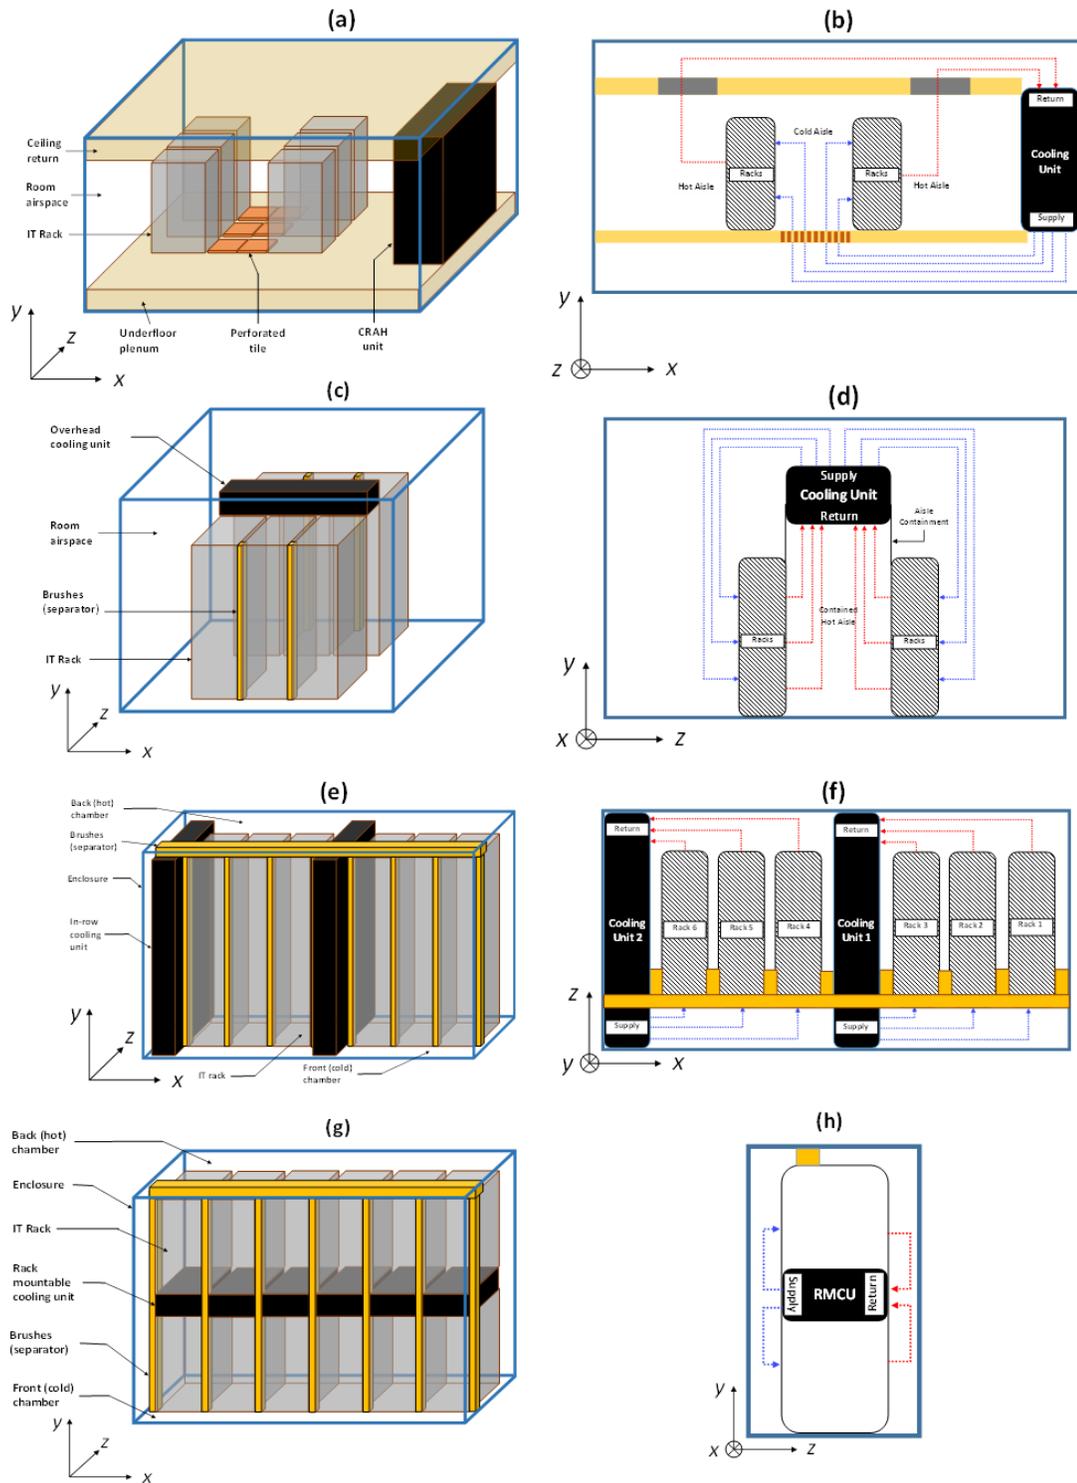


Figure 4.2: Schematic representation of the geometries used to simulate four cooling architectures: (a)-(b) UFAD, (c)-(d) OHAD, (e)-(f) IRC, (g)-(h) RMCU. Blue and red dotted lines depict cold and hot airflows through the racks, respectively.

Based on the sizes of the CRAH units available for the different configurations, e.g., the internal heat exchanger size, number, and capacity of fans, a cooling architecture requires a specific number of cooling units to handle the heat load of Table 4.1. For UFAD cooling (Figures 4.2a and 4.2b), a room-based cooling system is used that consists of a CRAH unit, underfloor plenum, perforated tiles, and a ceiling return to the CRAH. Cold air from the CRAH unit passing through the underfloor plenum is delivered into the cold aisle through perforated tiles. The perforated tiles are modeled through an additional momentum source using the modified body force (MBF) model [32, 33]. Racks are arranged in two rows, where each row consists of 3 racks. The server intakes are placed in the cold aisle, and hot air from the exhaust is guided to the ceiling return using two overhead ducts.

For OHAD cooling (Figures 4.2c and 4.2d), we consider a contained hot aisle [19, 20] and three overhead CRAH units, where each unit handles two racks. The cold air is delivered to the room from the overhead CRAH unit, whereupon it passes through the server, is exhausted into the hot aisle, and then drawn back into the CRAH unit.

For IRC cooling (Figures 4.2e and 4.2f), two CRAH units are employed, where each unit handles 3 IT racks. All the six racks and two cooling units are arranged in a single row within a container, as shown in Figure 4.2e. Cold air delivered by in-row cooling units passes through the servers and is released into the hot aisle (or back chamber).

The RMCU cooling scheme utilizes a rack-mountable CRAH unit of 2U size for each rack, resulting in 6 CRAH units for our case. These cooling units are placed at the midway height of the rack and deliver cold air vertically (Figures 4.2g and 4.2h).

The IT racks are represented by the recirculation boundary condition available in ANSYS Fluent that uses a thermodynamic energy balance to determine the rack exhaust temperature for a specified heat load and flowrates through servers. The flowrate and power consumption for the servers are known for a specific DC utilization. The CRAH unit inlet and exhaust zones are emulated using the inlet mass flow at a certain temperature (which is the cooling unit setpoint) and outlet pressure for a target mass flow rate, respectively. Table 4.2 shows the specifications of cooling units used for different architectures.

A steady-state assumption is imposed, which results in temporally invariant airflow through the racks and cooling units, the heat loads through the servers, and the air delivery temperature from the cooling units. In commercial DCs, the gaps between the racks or those between a rack and its containment walls produce mixing of the hot and cold airstreams, but this is usually minimized by inserting high-density air blocking brushes at these locations [6, 31]. The brushes provide flow resistance and minimize momentum and energy transport across aisles. To represent the effect of brushes in the CFD simulation, they are modeled as porous media using a power-law resistance so that [6],

$$\Delta p = -C_0 |v|^{C_1}, \quad (4.1)$$

where, $C_0=11$ and $C_1=1.15$ are determined from experiments.

Table 4.2: Geometric considerations for CFD simulation

Architecture	Maximum cooling capacity of each unit (kW)	Size of each cooling unit (m ³)	Maximum volume flow rate of air through each unit (m ³ s ⁻¹)	Number of cooling units used for the present study
IRC	33	1.2 × 0.3 × 2.0	2.01	2
UFAD	66	1.2 × 2.4 × 3.0	4.02	1
RMCU	11	0.6 × 0.1 × 1.0	0.67	6
OHAD	22	0.6 × 0.3 × 1.2	1.34	3

4.3.2 Exergy loss: server/rack component

The electrical power consumption of each server is assumed to be fully converted into heat, \dot{Q}_s , resulting in a server exhaust temperature,

$$T_o = T_i + \frac{\dot{Q}_s}{\dot{m}_s c_{pa}}, \quad (4.2)$$

where T_i denotes the inlet temperature to the server, and \dot{m}_s the mass flow of air through servers. For each server, the effective surface temperature T_s is calculated using a lumped thermal resistance approximation [17, 34],

$$T_s = T_i + \frac{\dot{Q}_s}{\dot{m}_s c_{pa} \left[\exp\left(-\frac{1}{\dot{m}_s c_{pa} R_s}\right) - 1 \right]}, \quad (4.3)$$

where R_s denotes the thermal resistance for the 1U server, which is 0.065 K/W [17].

With T_s , T_i and T_o known, the total exergy destruction with N servers $\dot{\psi}_{d,s}$ (see Figure 4.1),

$$\dot{\psi}_{d,s} = N \left\{ \dot{m}_s c_{pa} \left[(T_i - T_o) - T_o \ln \left(\frac{T_i}{T_o} \right) \right] + \dot{Q}_s \left[1 - \frac{T_o}{T_s} \right] \right\}. \quad (4.4)$$

We consider a DC with six racks, each rack having 40 1U servers with identical exergy destruction. Since the rack load is constant, the exergy destruction attributed to the server is also constant. The salient characteristics of the servers and IT racks are listed in Table 4.3.

Table 4.3: Server and IT rack parameter specifications.

Parameter	Value
Model	HP ProLiant DL360 G5
Server chassis size	1U
Power Consumption	262.5 W
Airflow rate	$0.01415 \text{ m}^3 \text{ s}^{-1}$
Thermal resistance	0.65 K W^{-1}
Number of servers	40
Rack type	Standard 42U racks (1U = 1.75 inch)
Number of racks	6

4.3.3 Exergy loss: airspace component

The exergy destruction associated with the airspace helps identify inefficiencies, cooling air maldistribution, and mismanagement of cooling architectures. This term is important for assessing the impact of different geometries, airflow management strategies, and cooling unit operating conditions on the exergy loss.

The airflow velocity, pressure, temperature, and turbulence in each control volume are determined through CFD simulations for different cooling architectures and operating parameters, as described in Table 4.2. The exergy destruction formulation for turbulent flows utilizing the $k - \epsilon$ model is as follows [18, 35-38].

$$\dot{S}_{VD}''' = \frac{\mu}{\bar{T}} \left\{ 2 \left[\left(\frac{\partial \bar{u}}{\partial x} \right)^2 + \left(\frac{\partial \bar{v}}{\partial y} \right)^2 + \left(\frac{\partial \bar{w}}{\partial z} \right)^2 \right] + \left(\frac{\partial \bar{u}}{\partial y} + \frac{\partial \bar{v}}{\partial x} \right)^2 + \left(\frac{\partial \bar{u}}{\partial z} + \frac{\partial \bar{w}}{\partial x} \right)^2 + \left(\frac{\partial \bar{v}}{\partial z} + \frac{\partial \bar{w}}{\partial y} \right)^2 \right\} + \frac{\rho \varepsilon}{\bar{T}}, \quad (4.5)$$

$$\dot{S}_{HT}''' = \frac{\lambda}{\bar{T}^2} \left(1 + \frac{C_\mu}{\alpha \text{Pr}_t} \frac{k^2}{\varepsilon} \right) \left[\left(\frac{\partial \bar{T}}{\partial x} \right)^2 + \left(\frac{\partial \bar{T}}{\partial y} \right)^2 + \left(\frac{\partial \bar{T}}{\partial z} \right)^2 \right], \quad (4.6)$$

$$\dot{\psi}_{HT}''' = T_0 \dot{S}_{HT}''', \quad (4.7)$$

$$\dot{\psi}_{VD}''' = T_0 \dot{S}_{VD}''', \quad (4.8)$$

$$\dot{\psi}_d''' = \dot{\psi}_{HT}''' + \dot{\psi}_{VD}''', \text{ and} \quad (4.9)$$

$$\dot{\psi}_{d,a} = \iiint \dot{\psi}_d''' dx dy dz. \quad (4.10)$$

Airspace exergy loss in a DC consists of two components, (a) exergy destruction due to unwanted heat transfer between the cold and hot air streams and (b) exergy destruction due to pressure drop or turbulent velocity fluctuations. The heat transfer is determined using Eq. 4.6, and the velocity gradient component is determined using Eq. 4.5. Once the CFD solution converges, Eqs. (4.5)-(4.9) are solved using a user-defined function (UDF) code in ANSYS Fluent 18.0 to obtain the exergy destruction inside each control volume. Subsequently, the volume integral of exergy destruction is determined using Eq. 4.10, which represents the loss of available energy.

4.3.4 Exergy loss: CRAH unit component

The air-water heat exchanger within the CRAH unit is the major exergy loss component. For the heat exchanger, the velocity gradient and turbulent exergy loss components have been shown to be insignificant [17, 18, 39]. Additional assumptions include (a) no

boundary heat transfer between the system and its surroundings and (b) negligible kinetic and potential energy changes. Thus, the exergy loss within the CRAH unit can be represented through the exergy balance across control volume (see Figure 4.1),

$$\dot{\psi}_a = \dot{m}_a C_{pa} \left[(T_{ha} - T_{ca}) - T_0 \ln \left(\frac{T_{ha}}{T_{ca}} \right) \right], \quad (4.11)$$

$$\dot{\psi}_w = \dot{m}_w C_{pw} \left[(T_{cw} - T_{hw}) - T_0 \ln \left(\frac{T_{cw}}{T_{hw}} \right) \right], \text{ and} \quad (4.12)$$

$$\dot{\psi}_{a,h} = \dot{\psi}_a + \dot{\psi}_w, \quad (4.13)$$

where $\dot{\psi}_a$ denotes the exergy lost by the airside and $\dot{\psi}_w$ that gained by the waterside.

From Eqs. (4.11)-(4.13), we note that the exergy loss associated with the heat exchanger is a function of the air mass flowrate \dot{m}_a , return temperature of air to CRAH T_{ha} , supply air temperature from CRAH T_{ca} , mass flowrate of water \dot{m}_w , inlet water temperature to CRAH T_{cw} , and outlet water temperature from CRAH T_{hw} . For a prescribed airflow and supply temperature from the CRAH, the return temperature to the CRAH unit is determined from the CFD simulations using the UDF in Fluent.

The airside parameters available from the CFD simulation are coupled with the waterside parameters using a code is written in MATLAB 2015 that models the heat exchangers within the CRAH. The code employs the $\varepsilon - NTU$ method [6, 17, 30] for coupling the waterside and airside parameters by specifying the type and size of the heat exchanger inside the CRAH. The dimensions and types of heat exchangers selected for the IRC, UFAD, RMCU, and OHAD cooling configurations are shown in Table 4.4. The water flowrates for different heat exchangers are based on commercially available CRAH units.

Table 4.4: Heat exchanger characteristics [19, 40, 41]

Architecture	Heat exchanger type	Heat exchanger size (m ³)	Maximum volume flow rate of water (m ³ s ⁻¹)	Number of heat exchangers used
IRC	Fin-tube	0.20 × 0.70 × 2.00	0.0009	2
UFAD	Fin-tube	0.20 × 1.50 × 2.25	0.0050	1
RMCU	Plate-fin	0.30 × 0.35 × 0.12	0.0004	6
OHAD	Fin-tube	0.30 × 1.20 × 0.60	0.0007	3

4.3.5 Exergy loss: chiller component

We consider an 88 kW chiller [26], which consists of a vapor compression refrigeration cycle and an ambient air-cooled condenser that supplies the required chilled water to the CRAH unit. Considering the chiller boundary as the control volume for energy interactions with the surroundings, the energy balance requires the electrical energy input into the chiller, energy transport through the chilled waterside, and energy interactions in the ambient airside. The chiller is assumed to be insulated. Thus the boundary heat transfer component is absent.

For the abovementioned conditions, the exergy destruction for the chiller (see Figure 4.1),

$$\begin{aligned} \dot{\psi}_{d,c} = & \dot{W} + \dot{m}_w c_{pw} \left[(T_{hw} - T_{cw}) - T_0 \ln \left(\frac{T_{hw}}{T_{cw}} \right) \right] \\ & + \dot{m}_f c_{pa} \left[(T_0 - T_{ha,0}) - T_0 \ln \left(\frac{T_0}{T_{ha,0}} \right) \right]. \end{aligned} \quad (4.14)$$

Imposing an energy balance across the chiller boundary control volume,

$$\dot{W} + \dot{m}_w C_{pw} (T_{hw} - T_{cw}) + \dot{m}_f C_{pa} (T_0 - T_{ha,0}). \quad (4.15)$$

Combining Eqs. (4.14) and (4.15), the chiller exergy loss,

$$\dot{\psi}_{d,c} = \dot{m}_w C_{pw} T_0 \ln \left(\frac{T_{cw}}{T_{hw}} \right) + \dot{m}_f C_{pa} T_0 \ln \left(\frac{T_{ha,0}}{T_0} \right). \quad (4.16)$$

To determine the exergy destruction associated with the chiller, the waterside, ambient airside, and the electric power consumption of the chiller must be determined. The total power consumption of the chiller is determined through the universal Ng-Gordon chiller model [42]. This model correlates the cooling load of the chiller evaporator, the air temperature entering the condenser, the desired setpoint of the chilled water leaving the evaporator, and the coefficient of performance (COP) as follows.

$$y = a_1 x_1 + a_2 x_2 + a_3 x_3, \quad (4.17)$$

$$x_1 = T_{cw} / Q_c, \quad (4.18)$$

$$x_2 = (T_0 - T_{cw}) / (T_0 \times Q_c), \quad (4.19)$$

$$x_3 = \left[\left\{ \left(\frac{1}{COP} \right) + 1 \right\} \times Q_c \right] / T_0, \text{ and} \quad (4.20)$$

$$y = \left\{ \left[\left(\frac{1}{COP} \right) + 1 \right] \times (T_{cw} / T_0) \right\} - 1. \quad (4.21)$$

The coefficient of performance, COP, is the ratio of chiller evaporator heat load, Q_c , to the total electrical power consumption by the chiller \dot{W} , and T_{cw} denotes the chilled water temperature leaving the evaporator. Air entering the condenser is assumed to have a temperature equal to T_0 . All temperatures are expressed in K and Q_c in kW. Data for Q_c , COP, T_{cw} , and T_0 are obtained from an 88kW chiller datasheet [26], Fitting that data to Eqs. (4.17)-(4.21) using a multivariate regression solver, the chiller model is,

$$y = 0.026x_1 + 28.872x_2 + 0.241x_3. \quad (4.22)$$

The waterside parameters (chilled water temperature supplied to CRAH, return water from CRAH, and water flowrate) are determined through the $\varepsilon - NTU$ method for the heat exchanger. Since the ambient air supply temperature is assumed to be equal to the thermodynamic dead state temperature (295.15 K), it is at zero exergy with respect to the ambient. The hot air temperature $T_{ha,0}$ returned to the ambient depends on the cooling load on the chiller and is determined from Eq. (4.15). The flow velocity of the chiller condenser fan \dot{m}_f is obtained from the datasheet of the chiller [26]. The characteristics of the chiller are provided in Table 4.5.

Table 4.5: Chiller parameter and specifications.

Parameter	Value
Model	TRANE CGAF C25
Rated maximum capacity	25 ton of refrigeration (88 kW)
Airflow of condenser fans	$10.22 \text{ m}^3 \text{ s}^{-1}$
Condenser air supply temperature	22 °C
Maximum chilled water flowrate	$0.0019 \text{ m}^3 \text{ s}^{-1}$
Maximum chilled water flowrate	$0.0057 \text{ m}^3 \text{ s}^{-1}$

4.4 Numerical procedure

The turbulent airflow field inside the DC is obtained by iteratively solving the RANS equations using the standard $k - \varepsilon$ model with ANSYS Fluent 18.0. The governing equations for mass, momentum, energy are solved for turbulence flow. Once convergence is obtained, the airspace exergy loss and the other components of exergy loss are calculated using thermodynamic energy-exergy balance equations. The simulation process is described in the appendix of Chapter 4.

4.4.1 Mesh independence study

A mesh independence analysis provides the minimum number of nodes required for accurately representing the flow physics. Four meshes with 10^6 , 1.7×10^6 , 2.3×10^6 , and 4×10^6 nodes provide different mesh geometries. The geometry with 2.3×10^6 nodes is selected for all cases due to its faster computation time while adequately determining the flow field. The RMSE for 60 different monitored nodes in the domain is lower than $1.5 \text{ }^\circ\text{C}$ for this grid [43]. Table 4.6 provides mesh independence results for the average rack inlet temperature, hot air return temperature to the cooling unit, and volume integral of the airspace exergy loss. These are the primary quantities of interest determined from the CFD simulations.

Table 4.6: Mesh independence results.

Number of cells	Average rack inlet temperature (K)	Return temperature of air to cooling unit (K)	Volume integral of airspace exergy loss (W)
1.0×10^6	295.2	310.7	480.3
1.7×10^6	295.9	311.3	495.7
2.3×10^6	296.5	311.9	510.1
4.0×10^6	296.5	311.9	510.1

4.4.2 Experimental validation

The experimental configuration consists of five operating racks and two in-row CRAH units that use chilled water for extracting heat from racks. Figure 4.3 provides a schematic representation of the experimental DC facility. The widths of the cold and hot aisles are 0.14 m and 0.20 m, respectively, which are separated from one another by brushes to prevent energy and momentum transport across them. The numerical model is for a steady-

state since the air temperature and power measurements demonstrate negligible temporal variations. The experimental operating conditions for racks and cooling units are provided in Tables 4.7 and 4.8, respectively.

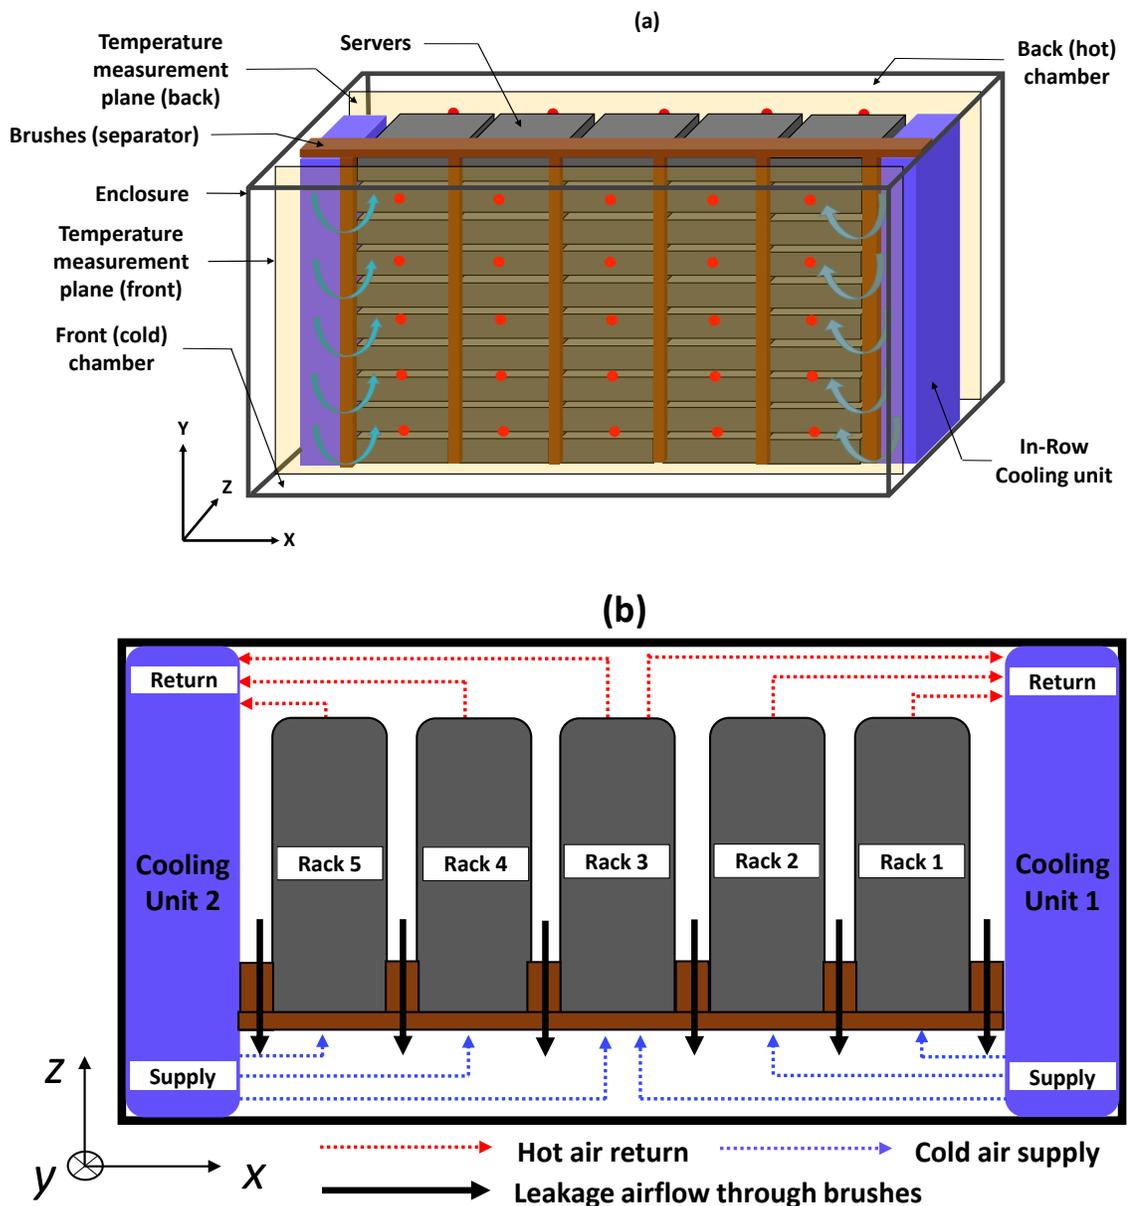


Figure 4.3: (a) Schematic representation of the experimental DC that has separated hot and cold chambers with 5 IT racks and 2 IRC units contained between them. In the aisle

chamber, cold air exits the IRC and is drawn through the servers. Hot air exits the servers in the back aisle and is drawn into the IRC. Leakage airflow occurs through the brushes (or separators) across the aisles. Red dots indicate positions identifying the locations of temperature probes that are placed along the halfway plane of the width of the cold ($Z = 0.06$ m) and hot ($Z = 1.22$ m) aisles. (b) top cross-sectional view showing salient airflows. The direction of leakage airflow depends on the pressure difference across the front and back chambers.

Table 4.7: Rack operating conditions.

Rack	Volume flow rate of air ($\text{m}^3 \text{s}^{-1}$)	It load (kW)
1	0.22	3.5
2	0.20	3.9
3	0.24	4.1
4	0.20	3.7
5	0.22	3.9

Table 4.8: Cooling unit operating conditions.

Case	Cooling unit	Volume flow rate of air ($\text{m}^3 \text{s}^{-1}$)	Cooling unit setpoint ($^{\circ}\text{C}$)
1	Left	0.41 + 0.41	18
	Right		
2	Left	0.51 + 0.51	
	Right		

The heat transfer and fluid flow are simulated following the methodology described earlier. The model assumes a complete conversion of the total IT load referred to in Table 4.7 into heat. The in-row cooling geometry shown in Figure 4.3 consists of a hot chamber, a cold chamber, and brushes that separate hot and cold chambers. Temperature measurements within the hot and cold chamber are performed with DS1820 temperature sensors that have an accuracy of ± 0.5 $^{\circ}\text{C}$ and are connected to an Arduino Mega development board that logs the data from 25 equidistant positions within each chamber. This provides the two-dimensional temperature distribution within each aisle along a plane

halfway through the depth of aisles at $Z = 0.06$ m for the cold aisle and $Z = 1.22$ m for the hot aisle and represented by the red dots in Figure 4.3a. The percentage of error between CFD and experiments relative to the experimentally measured temperatures at each location is defined as follows,

$$\Delta = \frac{T_{CFD} - T_{EXP}}{T_{EXP}} \times 100. \quad (4.23)$$

Figure 4.4 shows contours of the distribution of the relative percentage error Δ between the CFD simulations and experiments for the scenarios presented in Table 4.8. For both cases in the cold aisle, eight locations show a 9% deviation from the experiments, while the remainder show Δ values of 5% or lower. For the hot aisle, there are only four positions for case 2 where $\Delta = 8\%$, whereas 11 locations show a deviation of 7.5% or lower. The primary reasons for differences between the experiments and CFD simulations are as follows. First, the power and network cable bundles placed in the hot aisle are not considered in the CFD simulations. This simplification is the primary source of error for temperature predictions in the hot aisle. Second, when the airflow into the server suction is lowered, the probability of hot air recirculating to the cold aisle increases, which diminishes the accuracy of predicted temperatures. However, a 10% temperature prediction deviation is considered acceptable for operational DCs [18, 25, 31].

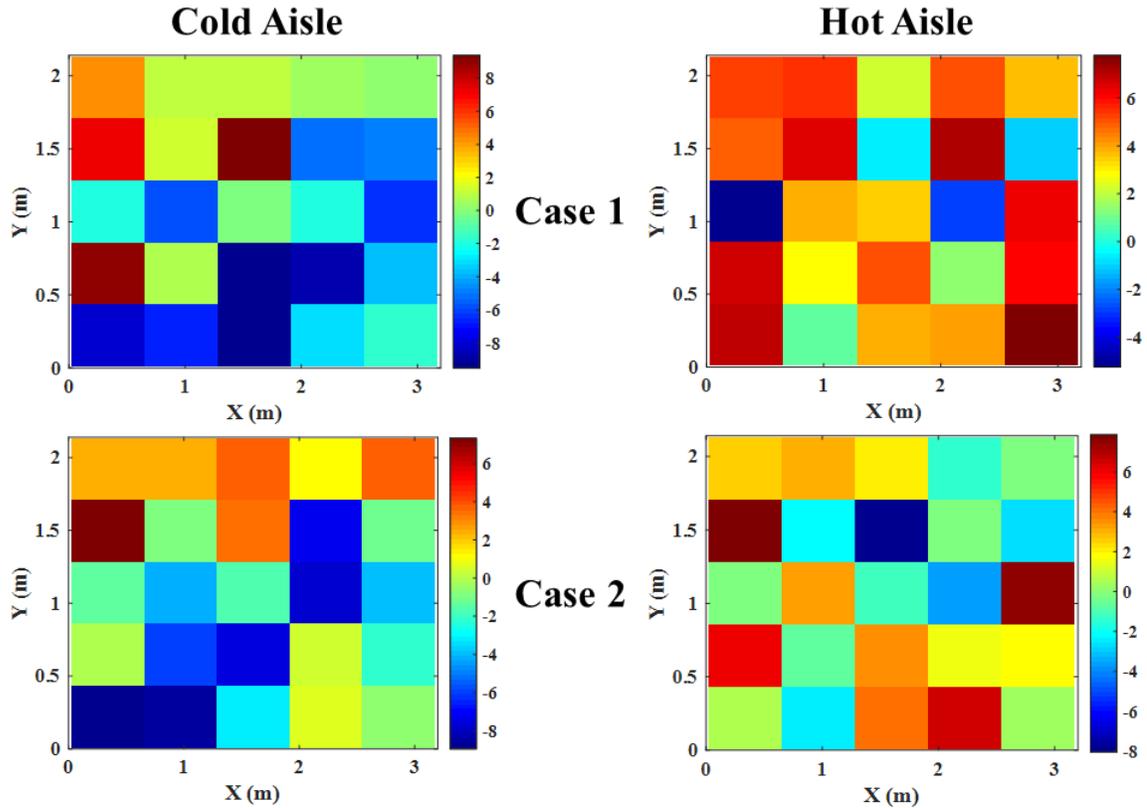


Figure 4.4: Distribution of percentage of relative error in predicted temperature Δ in cold and hot aisles for two different validation cases presented in Table 4.8.

4.5 Results and discussion

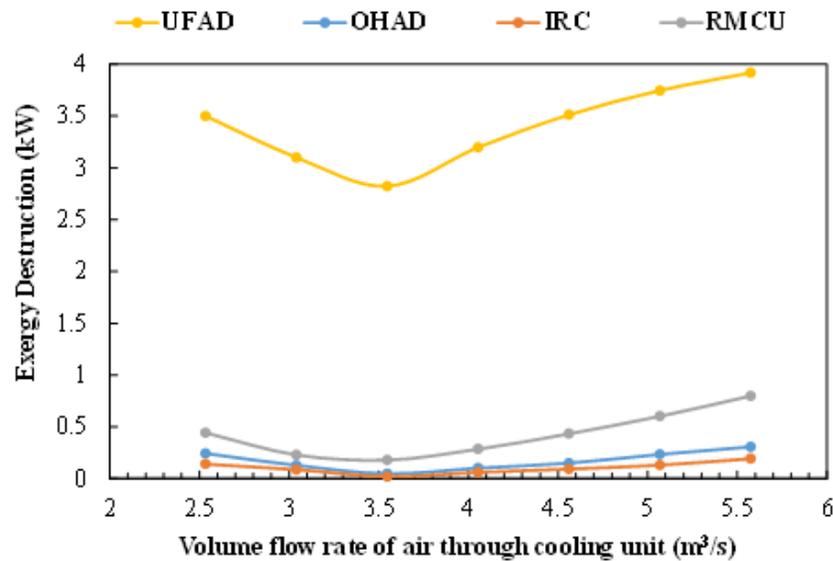
4.5.1 Airspace exergy loss

Airspace exergy destruction is caused by pressure drops, premature mixing of hot and cold airstreams, and turbulent fluctuations of energy and velocities. These effects are functions of (1) airflow path length, (2) cooling unit operating parameters, (3) number of cooling units, and (4) specific cooling architecture. The volume integral of the airspace exergy loss is determined for the scenarios and geometries depicted in Table 4.1.

Figure 4.5a shows the airspace exergy loss for different cooling architectures as a function of the CRAH unit airflow rate. The UFAD architecture has the highest exergy loss

(up to 4 kW), which is caused by several factors [8]. They include the (1) uncontained geometry that leads to higher recirculation and bypass, which in turn increase the exergy loss due to heat transfer, (2) use of perforated tiles that produce a pressure drop in the underfloor plenum, increasing exergy destruction due to velocity gradients, and (3) higher flow rates that cause the air to bypass the perforated tile closest to the cooling unit, forming a flow void inside certain regions in the underfloor plenum, resulting in air ingress back from the room to the underfloor. The heat transfer exergy loss contributes ~67% of the total exergy destruction, while 33% of overall loss is due to the velocity gradients [18].

Compared with the UFAD architecture, the other geometries (OHAD, IRC, and RMCU) show up to 80% percent lower overall exergy loss in the airspace. This exergetic improvement is made possible by the corresponding localized air delivery schemes that reduce recirculation and bypass while also eliminate pressure drop in the underfloor plenum [6, 18, 19].



(a)

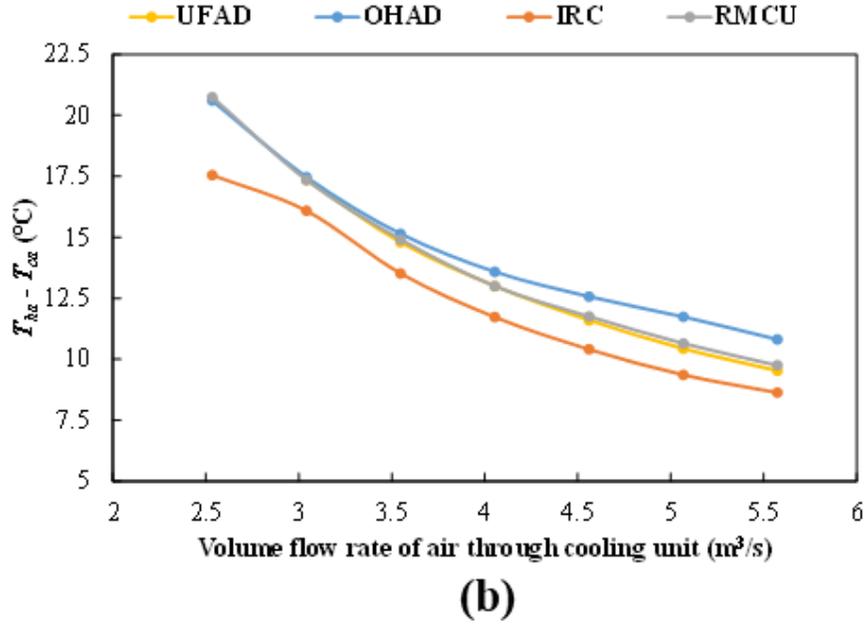


Figure 4.5: (a) airspace exergy loss, and (b) $T_{ha} - T_{ca}$ as a function of cooling unit airflow for different geometries.

For all geometries, increasing the airflow above a certain minimum ($\sim 3.5 \text{ m}^3 \text{ s}^{-1}$) increases the airspace exergy loss. This flowrate corresponding to the minimum exergy loss is the critical airflow for the DC, which occurs when the cooling unit airflow matches the total server air suction. Lower airflows lead to hot air recirculation occurs while higher values lead to server bypass. Hence, it is essential to identify the critical airflow when new servers are installed or the idle servers are turned off [6, 21, 23]. Figure 4.5b shows that increasing airflow monotonically decreases the difference between the supply and return temperatures ($T_{ha} - T_{ca}$). Since the airflow drawn by the servers is constant for all cases, the excess cold air is bypassed to the hot aisle through the porous brushes, lowering the mean air temperature in the hot aisle and the CRAH return.

Thus, several methods can be adopted to minimize airspace exergy loss, including (1) reducing the airflow path length by switching to a localized cooling architecture, (2)

implementing a contained air delivery scheme, (3) matching the total server suction with the cooling unit airflow using a predictive control algorithm, (4) decreasing the interaction between hot and cold aisles by using denser air-blocking brushes, and (5) hindering leakage across racks.

4.5.2 Heat exchanger exergy loss

Figure 4.6 presents the effect of cooling unit airflow variations on the CRAH exergy loss for different cooling architectures. The same fin-tube heat exchangers are used in the CRAH units for the UFAD and IRC architectures. In both cases, similar exergy loss variations are observed as the cooling unit airflow increases, where the loss first increases with increasing airflow and then decreases. However, the magnitude of exergy loss for the UFAD is 30% lower than for the IRC due to the higher effectiveness of the heat exchanger and the number of cooling units employed (one for UFAD and two for IRC). The OHAD architecture has nearly invariant exergy loss for all cases, but that for the RMCU progressively decreases by 33% as the volumetric flow rate is changed by 55%. We observe similar exergy destruction behaviors for UFAD and IRC heat exchangers and hence consider only one, the IRC, for additional comparison. However, the behaviors of RMCU and OHAD heat exchangers are very different, and both architectures must be considered. Thus, we examine the airside and waterside exergy destruction components of the IRC, RMCU, and OHAD architectures.

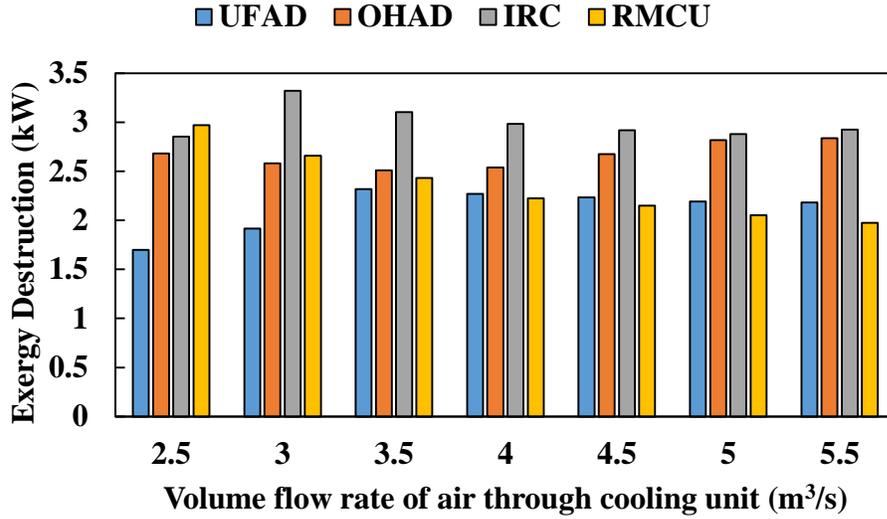


Figure 4.6: Heat exchanger exergy loss as a function of cooling unit airflow for different geometries.

The overall heat exchanger exergy loss is separated into (1) exergy gained by the waterside and (2) that lost by the air inside the heat exchanger. Figure 4.7a shows that exergy destruction for the OHAD heat exchanger has similar gradients on its waterside and airside, producing a nearly invariant total exergy loss with varying airflow. For the IRC architecture, the waterside exergy loss has a smaller gradient than on the airside (Figure 4.7b), but the overall exergy loss is again relatively invariant to changing airflow through the cooling unit. The RMCU heat exchanger is of plate-fin type, which has nearly invariant waterside exergy destruction with varying airflow (Figure 4.7c). However, due to the larger cold air bypass at elevated airflow, the return temperature to the heat exchanger decreases significantly (Figure 4.5b), decreasing $(T_{ha} - T_{ca})$ across it, resulting in a reduction of overall exergy loss.

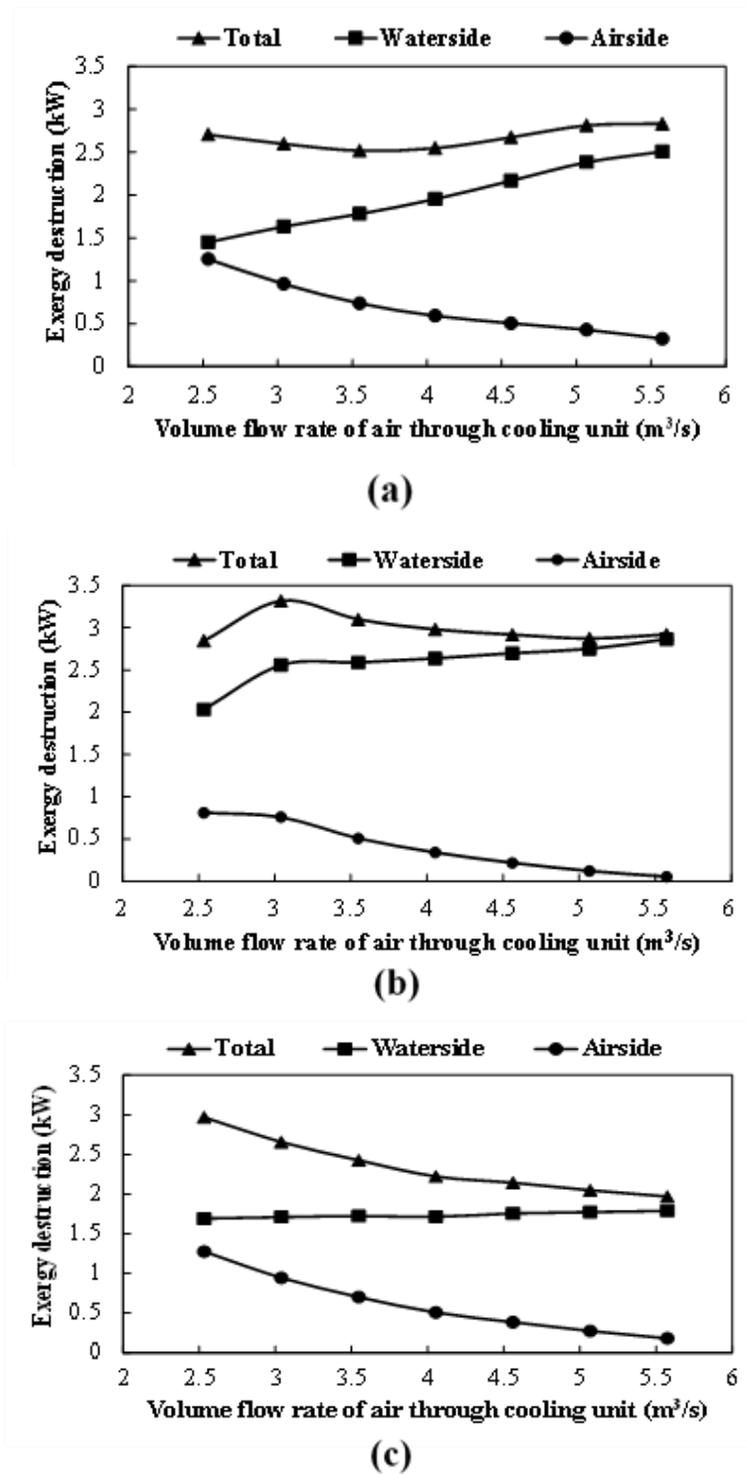


Figure 4.7: Airside and waterside components of the heat exchanger exergy loss for different geometries: (a) OHAD, (b) IRC, (c) RMCU

The exergetic efficiency for different heat exchangers [44],

$$\eta_{e,H} = \frac{\dot{\psi}_a}{\dot{\psi}_w} \times 100. \quad (4.24)$$

For the OHAD, IRC, RMCU, and UFAD architectures, the influence of cooling unit airflow on exergy efficiency is presented in Figure 4.8. When the airflow increases, the return temperature to the CRAH unit decreases due to the higher cold air bypass, which reduces the airside exergy destruction. This makes it more difficult to transfer heat from the airside to the waterside and reduces the exergetic performance of the heat exchanger. The fin-tube heat exchanger for the IRC has the lowest exergetic efficiency due to its lowest value of $(T_{ha} - T_{ca})$ as compared to the heat exchangers for RMCU and OHAD architectures (Figure 4.5b).

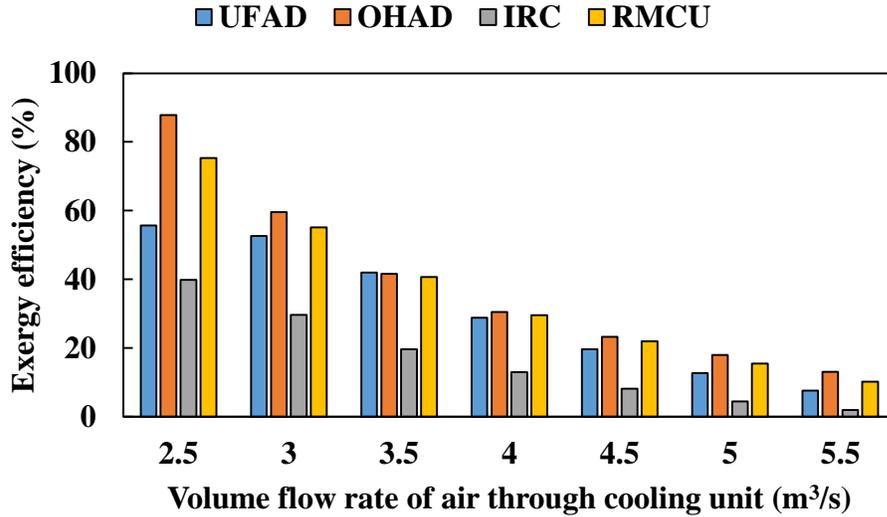


Figure 4.8: Influence of cooling unit airflow on the exergetic efficiencies of the heat exchangers used in the different architectures.

Although DCs typically employ additional cooling capacity in practice to support high network traffic loads, from an exergetic perspective, this increases exergy loss since the coolant temperature can have values below T_0 . Using a heat exchanger with greater

effectiveness reduces irreversibility, decreasing the additional provisioning for cooling that is required [39].

4.5.3 Chiller exergy loss

Figure 4.9 presents the variation in chiller exergy loss for the different architectures as the airflow through the cooling unit is increased. By increasing the cold airflow within a DC for a constant cold air delivery setpoint (see Table 4.1), the power consumption required for cooling also increases because this requires chilled water at lower temperatures. Hence, the temperature difference across the condenser side of the chiller increases since the chiller fans rotate at a constant speed and draw in air at a constant ambient temperature. This leads to an increment in the chiller exergy loss.

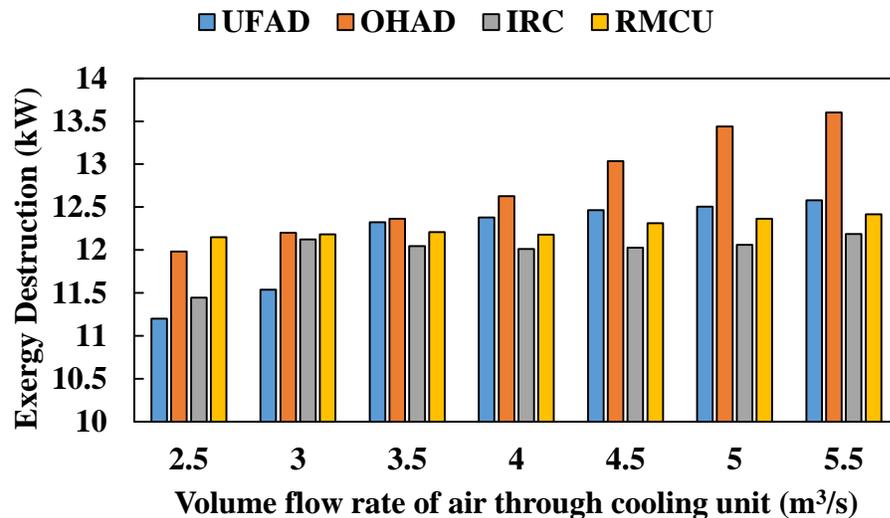


Figure 4.9: Influence of cooling unit airflow on the chiller exergy loss for different geometries.

The maximum exergy loss (~ 12.7 kW) inside the chiller occurs for the OHAD architecture when a $5.6 \text{ m}^3 \text{ s}^{-1}$ airflow is supplied through the cooling units. At higher

flowrates ($\sim 4.1\text{-}5.6 \text{ m}^3 \text{ s}^{-1}$), the OHAD architecture again provides the highest chiller exergy destruction followed by the UFAD, IRC, and RMCU systems. The chiller exergy loss for the RMCU is nearly invariant with changing airflow, while for the UFAD, OHAD, and IRC architectures, it increases due to the different characteristics of the heat exchangers that are used. The plate-fin heat exchanger used for the RMCU has a nearly invariant waterside exergy loss as the cooling unit airflow is increased (Figure 4.7c). Consequently, the changes in chiller exergy loss are not significant with increasing airflow.

Designers should be careful while implementing algorithms for fan speed control in response to temperature increments in the cold aisle since, in addition to fan power consumption, the chiller exergy loss increases significantly as the fans inside CRAH are rotated at faster speeds. For proper control that minimizes exergy loss, it is therefore essential to simultaneously monitor CRAH fans and chillers.

The chiller exergy destruction is made dimensionless to compare the chiller exergy efficiencies for different cooling configurations as follows [45].

$$\eta_{e,c} = \left\{ 1 - \frac{\dot{\psi}_{d,c}}{\dot{W} + \dot{m}_w C_{pw} \left[(T_{hw} - T_0) - T_0 \ln \left(\frac{T_{hw}}{T_0} \right) \right]} \right\} \times 100. \quad (4.25)$$

Figure 4.10 shows the exergy efficiency of the chiller, which depends on the exergy loss and the exergy input to the chiller, as a function of the volumetric airflow through the cooling units of four different geometries. Except for the RMCU architecture, the exergy efficiency of a chiller improves with increasing airflow through a cooling unit. For the RMCU, the chiller exergy efficiency is independent of airflow. The IRC has the highest

exergy chiller efficiency for the range of flowrates investigated. With increasing cooling unit airflow, the chiller heat load also increases and approaches its nominal design capacity. Near this heat load, the chiller works at higher COP and is, therefore, exergetically more efficient.

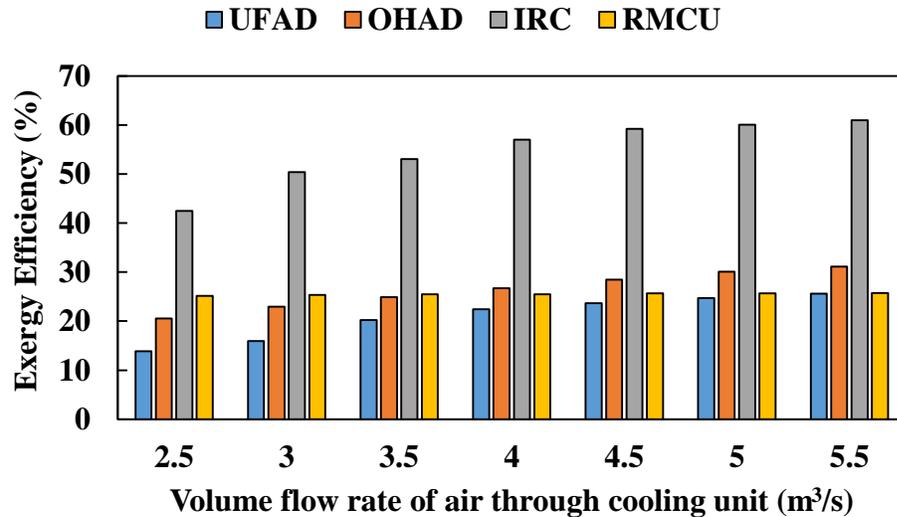
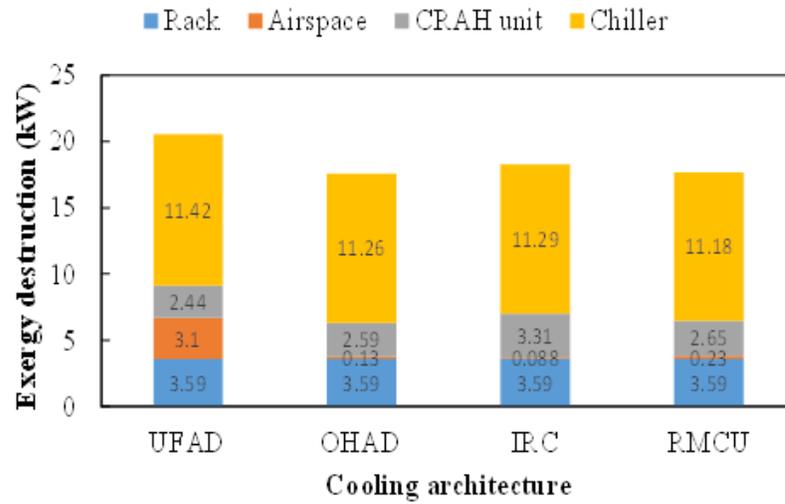


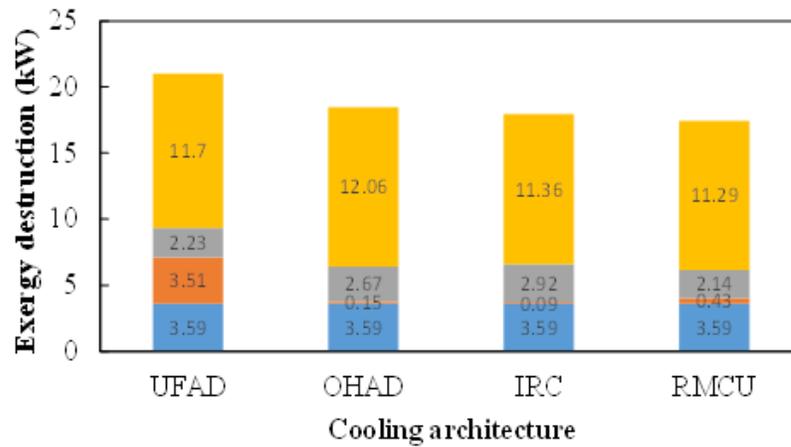
Figure 4.10: Influence of cooling unit airflow on the chiller exergy loss for different geometries.

4.5.4 Exergetic comparison of geometries under different scenarios

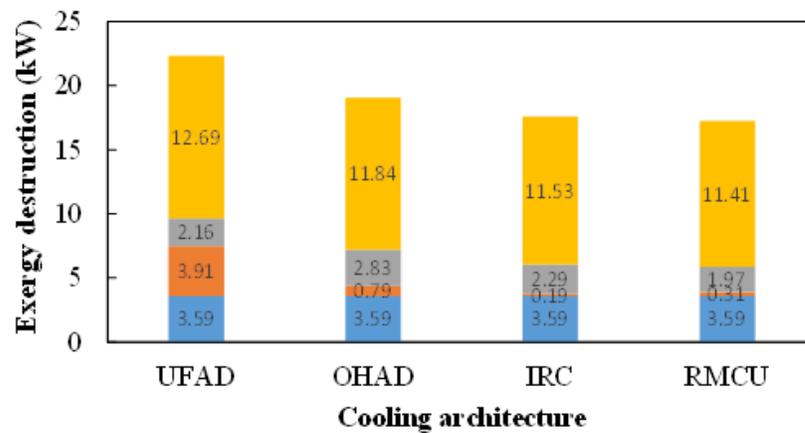
The exergy losses for the different components for the four architectures are compared in Figure 4.11 for three cooling unit flowrates, 3.04, 4.57, and 5.57 m³s⁻¹, and a specified cold air supply setpoint (see Table 4.1). The IRC and RMCU architectures outperform the room-based UFAD and OHAD legacy cooling by reducing hot air recirculation and cold air bypass. The rack exergy loss is the same for all cases due to the constant IT load imposed on the servers. The chiller exergy loss provides the highest contribution (~ 55% of the overall loss) to overall system inefficiency, followed by contributions from the rack and the heat exchanger.



(a)



(b)



(c)

Figure 4.11: Contributions of different exergy loss components for the different cooling architectures with three different CRAH flowrates: (a) 3.04, (b) 4.57, and (c) 5.57 m³ s⁻¹.

For the UFAD architecture, the airspace exergy loss is significant (~ 3-4 kW) as compared to that of other schemes. This is predominantly due to the uncontained UFAD geometry, which leads to premature cold and hot air mixing. For the RMCU, the single rack cooling produces the minimum exergy loss due to locally effective air delivery for all air flowrates investigated, as shown in Figure 4.11. Changing from the UFAD to the OHAD, IRC, or RMCU architectures can be represented by the parameter,

$$\Omega_g = \frac{\dot{\psi}_{UFAD} - \dot{\psi}_g}{\dot{\psi}_{UFAD}} \times 100 \quad (4.26)$$

Specifying the UFAD architecture as the base case, the overall exergy destruction for the OHAD, IRC, and RMCU architectures are compared in Table 4.9. For a 5.57 m³s⁻¹ airflow, the IRC and RMCU have 21.3% and 22.7% reductions in exergy loss, respectively, and the same trend holds for the 4.57 m³ s⁻¹ flowrate. The OHAD architecture provides nearly constant exergetic savings (~12% to 14%) for the three flowrates. As the flowrates are reduced to 3.04 m³ s⁻¹, the exergy loss is reduced to ~11-14%.

Table 4.9: Decrease in exergy loss for different architectures with respect to UFAD cooling.

Cooling unit airflow (m ³ s ⁻¹)	Ω_{UFAD}	Ω_{OHAD}	Ω_{IRC}	Ω_{RMCU}
5.57	----	14.8	21.3	22.7
4.57	----	12.2	14.6	17.0
3.04	----	14.5	11.1	14.1

Figure 4.12 provides a schematic of the exergy behaviors along with different components of an air-cooled DC, where arrow widths signify the relative magnitudes of

exergy transfers. Ambient air entering the condenser corresponds to zero exergy since the condenser inlet air temperature is assumed to be equal to the dead state temperature.

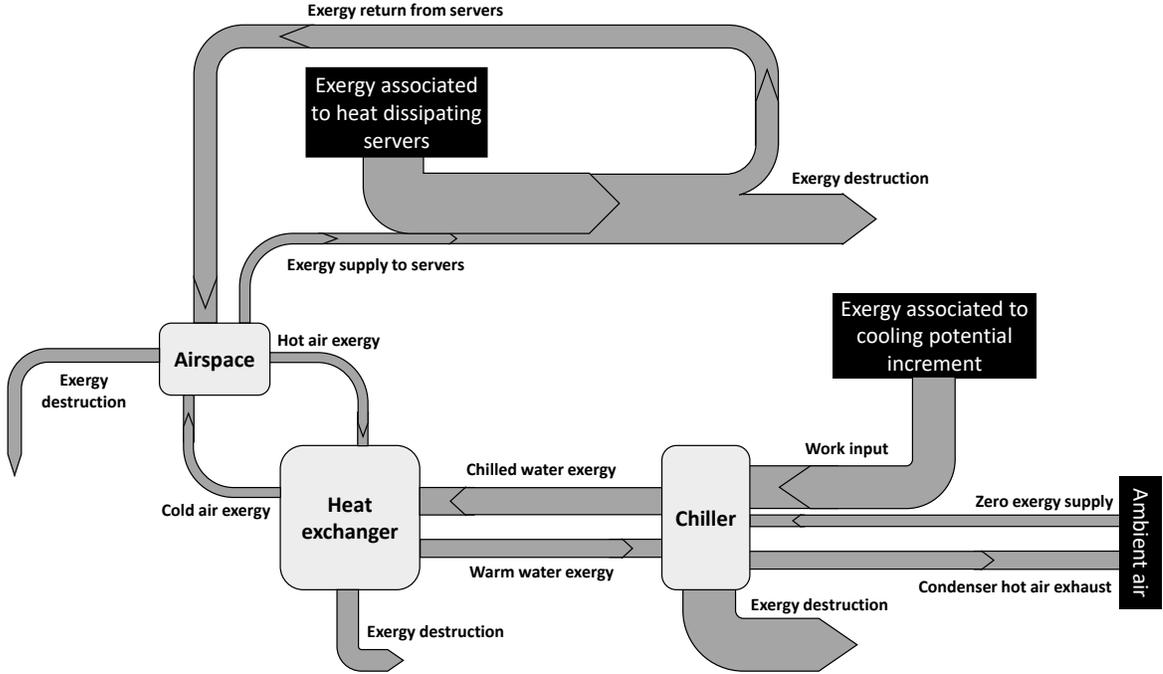


Figure 4.12: Schematic representation of exergy interactions in an air-cooled DC.

4.5.5 Dimensionless representation of exergy destruction

The overall exergy loss in the DC consists of four components, (1) in the racks due to heat dissipation, (2) in the airspace due to hot air recirculation and cold air bypass, (3) in the air-water heat exchanger, and (4) in the chiller. Because the cooling cycle exergy destruction is a function of total IT load \dot{Q}_{IT} on a DC, the exergy loss is made dimensionless with respect to it. A higher IT load leads to a significant increment in the overall exergy destruction. The dimensionless exergy destruction is expressed using Eq. (4.27).

$$\eta = \frac{\psi_{d,t}}{\dot{Q}_{IT}} = \frac{\psi_{d,s} + \psi_{d,a} + \psi_{d,h} + \psi_{d,c}}{\dot{Q}_{IT}} \quad (4.27)$$

The second dimensionless parameter is the Peclet number, which is the ratio of advective to the diffusive heat transport in a DC, where

$$Pe = \frac{U_c L_c}{\alpha}. \quad (4.28)$$

Here, α denotes the thermal diffusivity of air. The characteristic CRAH velocity U_c represents the effect of flow, and a characteristic length L_c accounts for the influence of geometry, where the length is taken as the maximum distance traveled by air within a specific architecture. Both U_c and L_c are unique to cooling architecture and CRAH airflow. Thus, two central phenomena, (1) hot air recirculation and (2) cold air bypass, are addressed through Pe . Minimizing the two airspace inefficiencies associated with flow and geometry reduces the cascading exergy losses that occur in the CRAH and chiller. For the UFAD and IRC architectures, L_c is measured horizontally since air travels in that direction, whereas for OHAD and RMCU, it is taken along the vertical direction of air travel.

The third ratio is the dimensionless dead-state temperature ratio [18, 35, 39],

$$\zeta = \frac{T_0}{T_{ha} - T_{ca}}. \quad (4.29)$$

which compares the ambient dead state temperature T_0 to the difference in hot and cold aisle temperatures $\Delta T = T_{ha} - T_{ca}$. Flow recirculation and bypass indirectly influence ΔT by either increasing or decreasing the return air temperature to CRAH. A higher value increases hot air recirculation, producing a greater temperature nonuniformity in the airspace, which increases exergy loss. Further, raising the dead state temperature increases heat transfer irreversibility between the system and its surroundings, influencing the overall

system irreversibility. Therefore, ζ is an essential parameter for comparing the enthalpy of the dead state the enthalpy rise across the hot and cold aisles in a DC.

Since the literature does not contain a specific dimensionless relation to express exergy loss as a function of the working parameters inside a DC, we correlate η with Pe and ζ . Figure 4.13 shows that the exergy loss per unit IT load η decreases with increasing $(Pe \times \zeta)$, i.e., increasing either Pe or ζ reduces exergy destruction in a DC, for which the best fit obtained from Figure 4.13 is,

$$\eta = 0.3[Pe \times \zeta]^{-0.1}. \quad (4.30)$$

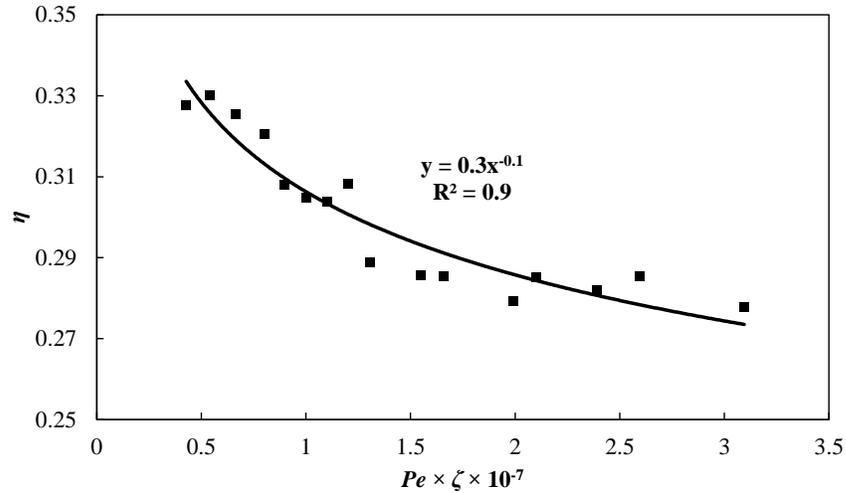


Figure 4.13: Dimensionless exergy loss η as a function of the dimensionless product $(Pe \times \zeta)$.

It is impractical to increase the air velocity in a cooling unit beyond a specific value, and hence geometric modifications must be considered to minimize exergy loss, which can involve changing the cooling architecture. Solely from an exergetic perspective, operating

a DC closer to ambient temperature decreases exergy destruction due to the reduction in exergy transfer to the ambient.

4.6 Conclusion

We provide an exergy destruction method to compare the performance of legacy DCs with emerging modular DCs. By employing the first and second laws of thermodynamics, system inefficiencies are represented as energy and exergy losses in a manner that can be used by DC designers and operators. Simulated airflows inform the thermodynamic energy and exergy balances used to determine the exergy loss in the different components of an air-cooled DC. The effects of cooling geometry and operating parameters on the exergy loss in these components are investigated. Salient findings include:

- The airspace exergy loss constitutes up to 20% of the total exergy loss for the UFAD architecture due to premature hot and cold air mixing and an underfloor pressure drop. For the other three architectures, the airspace contributes a much smaller ~5% of the total system exergy loss. Thereby, for DCs with UFAD architecture, the recirculation and bypass must be reduced by dynamically regulating airflows in real-time.
- The chiller exergy loss has the highest contribution (~ 55 to 60%) towards the total exergy loss for all geometries and scenarios investigated. We have considered a vapor compression refrigeration chiller, which is widely used in DCs. This exergy loss can be minimized by recovering waste heat if thermally driven chillers, i.e., absorption chillers and adsorption chillers, are used instead.
- Increasing the volumetric airflow of the cooling unit increases the chiller exergy loss progressively for all air delivery schemes. Hence, increasing the fan speed not only

increases fan power consumption but also leads to higher chiller power consumption, which increases the exergy loss in the chiller. Thus, an exergy-aware coordinated control of CRAH fans and chillers should be implemented to minimize exergy and energy losses simultaneously.

- For high cooling unit airflow, the exergy efficiency of the heat exchangers inside the CRAH decreases, whereas the chiller becomes exergetically efficient.
- The UFAD configuration produces the highest exergy loss, while DCs with the modular RMCU has the lowest exergy destruction among the four architectures. The RMCU has up to 23% lower overall exergy loss than the UFAD scheme.
- The dimensionless exergy destruction η decreases monotonically with increasing $(Pe \times \zeta)$.

Overall, we demonstrate the potential of second law analysis to improve air-cooled DC design, particularly to optimize its cooling and operation while reducing energy costs.

4.7 Appendix: Process for exergy loss calculation

The RANS equations are used in combination with the widely used $k - \varepsilon$ model to represent turbulence, and to resolve the characteristic airflow in DC. Governing equations for mass, momentum, energy, and turbulence, i.e., Eqs. (4.1A) through (4.8A) are numerically solved for each geometry and parametric test case using the commercial ANSYS Fluent 18.0 software. For a steady-state assumption adopted in this study, the continuity and RANS equations are,

$$\bar{\nabla} \cdot \bar{U} = 0 \tag{4.1A}$$

$$\bar{\nabla} \cdot (\bar{U}\bar{U}) = \frac{1}{\rho} \bar{\nabla} \cdot (\bar{\sigma} - \rho \bar{U}'\bar{U}') + \frac{1}{\rho} \bar{S} \quad (4.2A)$$

$$\bar{\sigma} = -PI + \mu\{\bar{\nabla}(\bar{U}) + [\bar{\nabla}(\bar{U})]^T\} \quad (4.3A)$$

Where, $\bar{\sigma}$ is the Newtonian stress tensor, μ is the dynamic viscosity of air, ρ is the density of air defined using an ideal gas equation of state, \bar{U} and \bar{U}' are the mean and turbulent velocity vectors, respectively, P is the pressure and I is the unit tensor. The momentum source term \bar{S} of the RANS equation is only valid for the zones that contain the air blocking brushes in DC. The source term is calculated in FLUENT using a power-law approximation model as described using Eq. (4.1). The governing equations for the utilized standard $k - \varepsilon$ turbulence model are as follows,

$$\bar{\nabla} \cdot (k\bar{U}) = \frac{1}{\rho} \bar{\nabla} \cdot \left(\frac{\mu_t}{\rho_k} \bar{\nabla}(k) \right) + \frac{2\mu_t}{\rho} S_{ij} \cdot S_{ij} - \varepsilon \quad (4.4A)$$

$$\bar{\nabla} \cdot (\varepsilon\bar{U}) = \frac{1}{\rho} \bar{\nabla} \cdot \left(\frac{\mu_t}{\rho_\varepsilon} \bar{\nabla}(\varepsilon) \right) + C_{1\varepsilon} \frac{\varepsilon}{k\rho} 2\mu_t S_{ij} \cdot S_{ij} - C_{2\varepsilon} \frac{\varepsilon^2}{k} \quad (4.5A)$$

$$\mu_t = \rho C_\mu \frac{k^2}{\varepsilon} \quad (4.6A)$$

where, k is the turbulent kinetic energy, ε is the viscous dissipation rate, S_{ij} is the deformation tensor, and μ_t is the turbulent viscosity represented using Eq. A6. The model constants for the standard $k - \varepsilon$ model are, $C_\mu = 0.09$, $\rho_k = 1.0$, $\rho_\varepsilon = 1.3$, $C_{1\varepsilon} = 1.44$, and $C_{2\varepsilon} = 1.92$.

The energy equation utilizing $k - \varepsilon$ model can be written as,

$$\bar{\nabla} \cdot (T\bar{U}) = \bar{\nabla} \cdot \left(\left[\frac{v}{Pr} + \frac{v_t}{Pr_t} \right] \bar{\nabla}(T) \right), \quad (4.7A)$$

$$Pr = \frac{\nu}{\alpha} \text{ and } \alpha = \frac{\lambda}{\rho C_p}, \quad (4.8A)$$

where, Pr is the Prandtl number represented using Eq. 4.8A, Pr_t is the turbulent Prandtl number, λ is the thermal conductivity of air, and C_p is the specific heat of air. For standard $k - \varepsilon$ model the value of Pr_t is 0.85.

The momentum and energy equations are discretized within the solution domain using a second-order upwind scheme, and first-order upwind is used to discretize the turbulent equations. The SIMPLE algorithm is used for pressure-velocity coupling. Exergy destruction within the airspace is calculated using Eqs. (4.5) – (4.9). Once the solver meets a specified convergence criterion when the residuals of all field variables reach 10^{-5} , the local exergy destruction within each control volume and subsequently its volume integral within the whole domain is computed using UDF in Fluent.

Once airspace convergence is obtained, the *ad hoc* heat exchanger model implemented in MATLAB determines the airside and waterside parameters required to calculate the exergy loss for the specific heat exchanger type (see Table 4.4). Further, the chiller power consumption and exergy destruction models are also used to determine the exergy loss in the chiller. A schematic representation of the numerical methodology is provided in Figure 4.14.

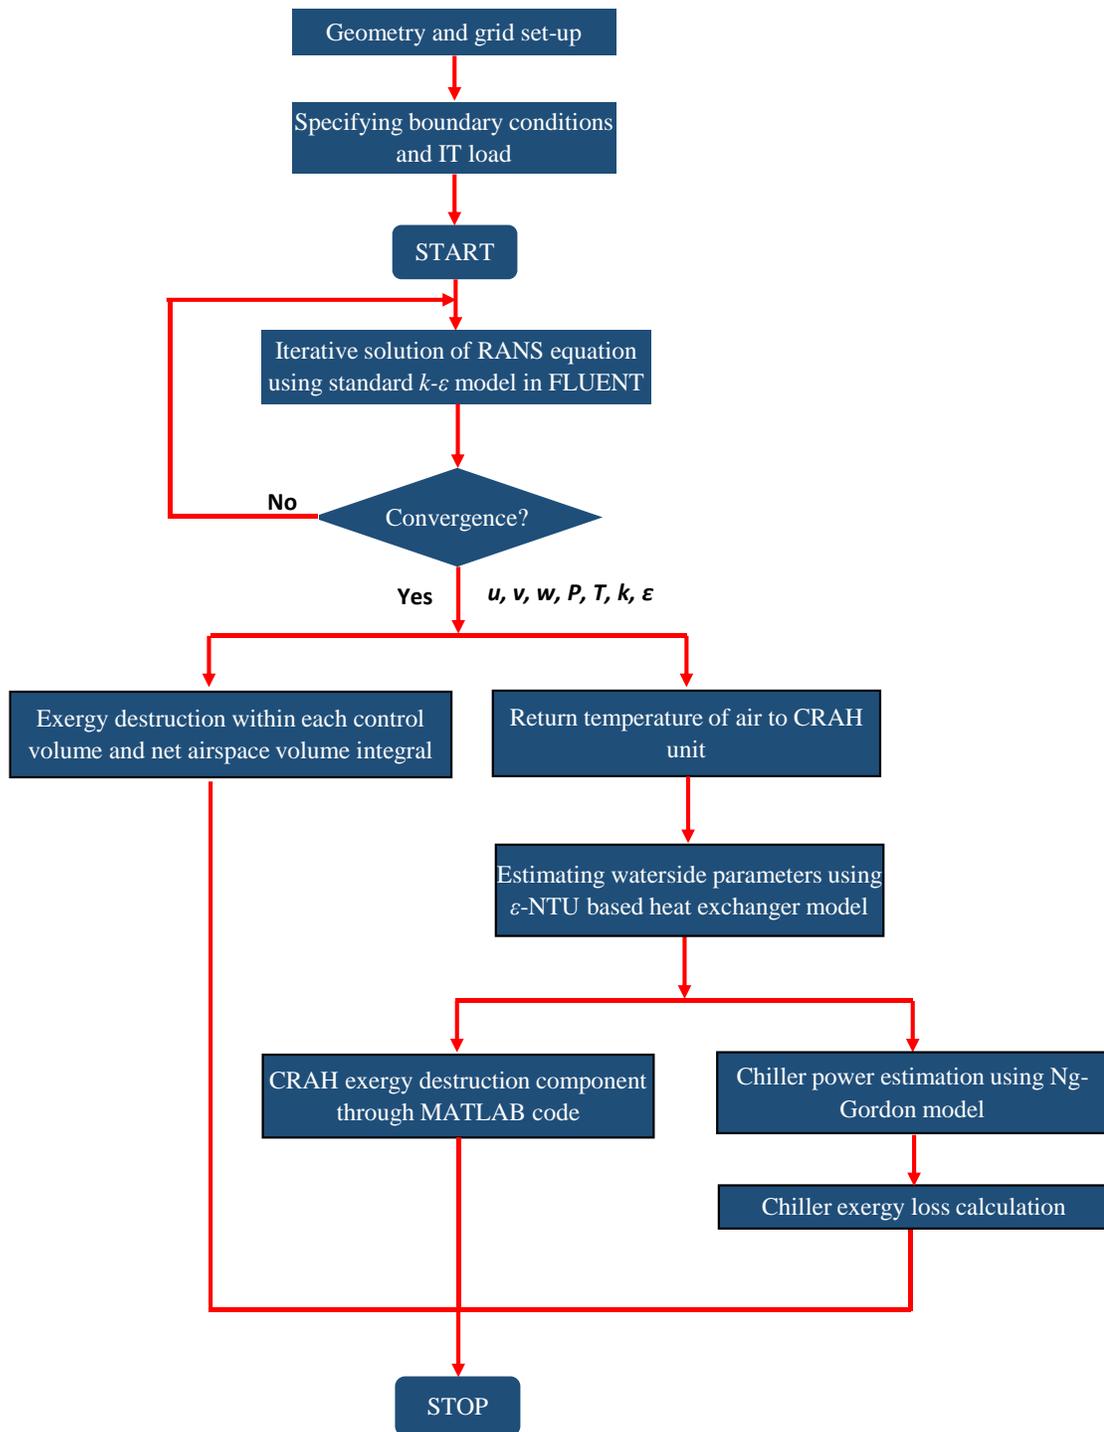


Figure 4.14: Schematic representation of the numerical procedure for computing exergy loss in the airspace, heat exchanger and chillers.

4.8 References

- [1] N. Rasmussen, "Air distribution architecture options for mission critical facilities," *Elektron Journal-South African Institute of Electrical Engineers*, vol. 22, p. 68, 2005.
- [2] A. H. Khalaj, T. Scherer, and S. K. Halgamuge, "Energy, environmental and economical saving potential of data centers with various economizers across Australia," *Applied energy*, vol. 183, pp. 1528-1549, 2016.
- [3] K. Ebrahimi, G. F. Jones, and A. S. Fleischer, "Thermo-economic analysis of steady state waste heat recovery in data centers using absorption refrigeration," *Applied Energy*, vol. 139, pp. 384-397, 2015.
- [4] S. Zimmermann, M. K. Tiwari, I. Meijer, S. Paredes, B. Michel, and D. Poulikakos, "Hot water cooled electronics: exergy analysis and waste heat reuse feasibility," *International Journal of Heat and Mass Transfer*, vol. 55, pp. 6391-6399, 2012.
- [5] S. Zimmermann, I. Meijer, M. K. Tiwari, S. Paredes, B. Michel, and D. Poulikakos, "Aquasar: A hot water cooled data center with direct energy reuse," *Energy*, vol. 43, pp. 237-245, 2012.
- [6] H. Moazamigoodarzi, P. J. Tsai, S. Pal, S. Ghosh, and I. K. Puri, "Influence of cooling architecture on data center power consumption," *Energy*, vol. 183, pp. 525-535, 2019.
- [7] K. Fouladi, J. Schaadt, and A. P. Wemhoff, "A novel approach to the data center hybrid cooling design with containment," *Numerical Heat Transfer, Part A: Applications*, vol. 71, pp. 477-487, 2017.

- [8] M. Xie, J. Wang, and J. Liu, "Evaluation metrics of thermal management in data centers based on exergy analysis," *Applied Thermal Engineering*, vol. 147, pp. 1083-1095, 2019.
- [9] K. Zhang, Y. Zhang, J. Liu, and X. Niu, "Recent advancements on thermal management and evaluation for data centers," *Applied thermal engineering*, vol. 142, pp. 215-231, 2018.
- [10] A. Capozzoli, G. Serale, L. Liuzzo, and M. Chinnici, "Thermal metrics for data centers: A critical review," *Energy Procedia*, vol. 62, pp. 391-400, 2014.
- [11] A. Beitelmal and D. Fabris, "Servers and data centers energy performance metrics," *Energy and buildings*, vol. 80, pp. 562-569, 2014.
- [12] A. J. Díaz, R. Cáceres, J. M. Cardemil, and L. Silva-Llanca, "Energy and exergy assessment in a perimeter cooled data center: The value of second law efficiency," *Applied Thermal Engineering*, vol. 124, pp. 820-830, 2017.
- [13] R. Khalid, A. P. Wemhoff, and Y. Joshi, "Energy and exergy analysis of modular data centers," *IEEE Transactions on Components, Packaging and Manufacturing Technology*, vol. 7, pp. 1440-1452, 2017.
- [14] X. Qian, Z. Li, and Z. Li, "Entransy and exergy analyses of airflow organization in data centers," *International Journal of Heat and Mass Transfer*, vol. 81, pp. 252-259, 2015.
- [15] A. J. Shah, V. P. Carey, C. E. Bash, and C. D. Patel, "Exergy analysis of data center thermal management systems," *Journal of Heat Transfer*, vol. 130, 2008.

- [16] A. Bhalerao, K. Fouladi, L. Silva-Llanca, and A. P. Wemhoff, "Rapid prediction of exergy destruction in data centers due to airflow mixing," *Numerical Heat Transfer, Part A: Applications*, vol. 70, pp. 48-63, 2016.
- [17] K. Fouladi, A. P. Wemhoff, L. Silva-Llanca, K. Abbasi, and A. Ortega, "Optimization of data center cooling efficiency using reduced order flow modeling within a flow network modeling approach," *Applied Thermal Engineering*, vol. 124, pp. 929-939, 2017.
- [18] L. Silva-Llanca, A. Ortega, K. Fouladi, M. del Valle, and V. Sundaralingam, "Determining wasted energy in the airside of a perimeter-cooled data center via direct computation of the Exergy Destruction," *Applied Energy*, vol. 213, pp. 235-246, 2018.
- [19] L. Silva-Llanca, M. Del Valle, A. Ortega, and A. J. Díaz, "Cooling effectiveness of a data center room under overhead airflow via entropy generation assessment in transient scenarios," *Entropy*, vol. 21, p. 98, 2019.
- [20] L. Silva-Llanca, M. del Valle, and A. Ortega, "The effectiveness of data center overhead cooling in steady and transient scenarios: comparison of downward flow to a cold aisle versus upward flow from a hot aisle," in *ASME 2015 International Technical Conference and Exhibition on Packaging and Integration of Electronic and Photonic Microsystems collocated with the ASME 2015 13th International Conference on Nanochannels, Microchannels, and Minichannels*, 2015.

- [21] H. Moazamigoodarzi, S. Pal, S. Ghosh, and I. K. Puri, "Real-time temperature predictions in it server enclosures," *International Journal of Heat and Mass Transfer*, vol. 127, pp. 890-900, 2018.
- [22] S. Pal, H. Moazami-goodarzi, S. Ghosh, K. Thorn, and I. K. Puri, "Plate-fin heat exchanger suitable for rack-mountable cooling unit," ed: Google Patents, 2019.
- [23] H. Moazamigoodarzi, S. Pal, D. Down, M. Esmalifalak, and I. K. Puri, "Performance of a rack mountable cooling unit in an IT server enclosure," *Thermal Science and Engineering Progress*, vol. 17, p. 100395, 2020.
- [24] J. Priyadumkol and C. Kittichaikarn, "Application of the combined air-conditioning systems for energy conservation in data center," *Energy and buildings*, vol. 68, pp. 580-586, 2014.
- [25] H. Moazamigoodarzi, R. Gupta, S. Pal, P. J. Tsai, S. Ghosh, and I. K. Puri, "Modeling temperature distribution and power consumption in IT server enclosures with row-based cooling architectures," *Applied Energy*, vol. 261, p. 114355, 2020.
- [26] TRANE, "Air-cooled liquid chillers, 10 to 60 Tons," ed, 2004.
- [27] ebm-papst, "EC centrifugal fan," ed, 2012.
- [28] A. S. o. Heating, Refrigerating, A.-C. Engineers, and A. N. S. Institute, *Thermal environmental conditions for human occupancy* vol. 55: American Society of Heating, Refrigerating and Air-Conditioning Engineers, 2004.
- [29] R. Steinbrecher and R. Schmidt, "Data center environments," *ASHRAE J*, vol. 42, p. e49, 2011.
- [30] W. M. Kays and A. L. London, "Compact heat exchangers," 1984.

- [31] A. Almoli, A. Thompson, N. Kapur, J. Summers, H. Thompson, and G. Hannah, "Computational fluid dynamic investigation of liquid rack cooling in data centres," *Applied energy*, vol. 89, pp. 150-155, 2012.
- [32] V. K. Arghode and Y. Joshi, "Modeling strategies for air flow through perforated tiles in a data center," *IEEE Transactions on Components, Packaging and Manufacturing Technology*, vol. 3, pp. 800-810, 2013.
- [33] V. K. Arghode, P. Kumar, Y. Joshi, T. Weiss, and G. Meyer, "Rack level modeling of air flow through perforated tile in a data center," *Journal of Electronic Packaging*, vol. 135, 2013.
- [34] D. W. Demetriou, H. S. Erden, H. E. Khalifa, and R. R. Schmidt, "Development of an IT equipment lumped capacitance parameter database for transient data center simulations," in *Fourteenth Intersociety Conference on Thermal and Thermomechanical Phenomena in Electronic Systems (ITherm)*, 2014, pp. 1330-1337.
- [35] A. Bejan, "Entropy generation minimization: The new thermodynamics of finite-size devices and finite-time processes," *Journal of Applied Physics*, vol. 79, pp. 1191-1218, 1996.
- [36] F. Kock and H. Herwig, "Entropy production calculation for turbulent shear flows and their implementation in CFD codes," *International Journal of Heat and Fluid Flow*, vol. 26, pp. 672-680, 2005.

- [37] F. Kock and H. Herwig, "Local entropy production in turbulent shear flows: a high-Reynolds number model with wall functions," *International journal of heat and mass transfer*, vol. 47, pp. 2205-2215, 2004.
- [38] H. Herwig and F. Kock, "Direct and indirect methods of calculating entropy generation rates in turbulent convective heat transfer problems," *Heat and mass transfer*, vol. 43, pp. 207-215, 2007.
- [39] A. Bejan, *Entropy generation minimization: the method of thermodynamic optimization of finite-size systems and finite-time processes*: CRC press, 2013.
- [40] RITTAL, "CRAC Precision climate control units for data centers, Model number: 3300.387," ed, 2014.
- [41] RITTAL, "C.W. TopTherm LCP rack/inline, Model number: 3311.130," 2015.
- [42] J. M. Gordon and K. C. Ng, *Cool thermodynamics*: Viva Books, 2008.
- [43] E. Cruz and Y. Joshi, "Inviscid and viscous numerical models compared to experimental data in a small data center test cell," *Journal of Electronic Packaging*, vol. 135, 2013.
- [44] T. K. Ray, A. Datta, A. Gupta, and R. Ganguly, "Exergy-based performance analysis for proper O&M decisions in a steam power plant," *Energy Conversion and Management*, vol. 51, pp. 1333-1344, 2010.
- [45] J. U. Ahamed, R. Saidur, and H. H. Masjuki, "A review on exergy analysis of vapor compression refrigeration system," *Renewable and Sustainable Energy Reviews*, vol. 15, pp. 1593-1600, 2011.

5 Energy and exergy based workload management for single-rack data centers

This chapter is reproduced from *Workload management for air-cooled data centers: An energy and exergy based approach*, *Energy* 209 (2020): 118484, **Rohit Gupta**, Hosein Moazamigoodarzi, SeyedMorteza MirhoseiniNejad, Douglas G. Down, and Ishwar K. Puri, <https://doi.org/10.1016/j.energy.2020.118485>, Copyright © 2021, with permission from Elsevier. The author of this thesis is the first author and the primary contributor to this publication.

5.1 Abstract

The energy required to cool an air-cooled data center (DC) contributes significantly to the cost of operation, which is further exacerbated due to a poor choice of cooling architecture and ineffective IT workload management. Although existing algorithms reduce energy consumption, they do not minimize thermodynamic irreversibility by design. We provide a tradeoff approach that simultaneously minimizes power usage effectiveness *PUE* and maximizes the exergy efficiency η_{2nd} . The temperature field is predicted inside a contained single-rack DC that is equipped with a rack-mountable cooling unit (RMCU) based on a mechanical resistance model for the fluid flow. This thermal model informs a multi-objective optimization framework based on a genetic algorithm to determine the optimal decision variables and tradeoffs for *PUE* and η_{2nd} . We investigate the interrelated effects of (1) guidelines that ensure the reliability of the IT equipment, (2) overall network traffic load, (3) spatial IT load distribution, (4) changes in cooling system variables, and (5) multi-

objective optimization. Results for the single rack system are presented in a scalable dimensionless form that is applicable for a multi-rack DC containing RMCUs. By considering the first and second laws of thermodynamics, this novel approach improves workload scheduling from both energy and exergy perspectives.

5.2 Introduction

Rapidly increasing computing demand has resulted in data centers (DCs) that are hyper-scale cloud facilities running at several petaFLOPS (floating-point operations per second), which now account for nearly 2-2.5% of the world's electricity demand [1], making them environmentally unsustainable and fiscally expensive. From chip to the chiller, a variety of thermal management and energy minimization strategies are employed to more effectively use the cooling energy consumed by a DC [2]. At the chip level, the thermal stress that degrades the chip and its energy consumption is managed through approaches such as dynamic voltage and frequency scaling (DVFS) [3], thread migration [4], frequency capping [5], and dynamic power management (DPM) [6]. At the server level, the most widely used strategy is to dynamically turn off unutilized servers, e.g., through server consolidation [7]. At the rack level, strategies to decrease cooling energy consumption include shortened air paths [8], intelligent workload assignment to servers [9], and switching to liquid-cooled servers [10]. At the facility level, overall energy reduction is typically based on renewable energy sources and harvesting waste DC heat [11].

Workload scheduling is a key tool for reducing DC operating costs while assuring IT equipment reliability. Algorithms assign the IT load to specific servers based on predefined criteria while adhering to thermal reliability guidelines. A strategy should

consider (1) cost, (2) computing efficiency, and (3) resilience. Workload scheduling and cooling system control can also use chip temperature-aware approaches [12, 13], where incoming compute loads are assigned to servers with the lowest chip temperatures, maintaining CPU temperatures below 75 °C, which also somewhat reduces the energy budget of a DC.

Another approach is to maintain a uniform temperature at the server exhaust through workload scheduling [14, 15], but this does not consider the cooling energy consumption *a priori*. The workload can be assigned based on the server inlet temperature and the IT load of neighboring servers [9] or to minimize hot air recirculation inside the DC [16]. These approaches consider a static heat-recirculation matrix based on the pressure-flow characteristics obtained from a computational fluid dynamics (CFD) simulation but often lead to erroneous dynamic temperature predictions. Since thermal-aware methods attempt to optimize energy through indirect strategies that reduce overcooling, they do not directly consider energy consumption in their objective functions.

Energy consumption can also be placed within the objective function of an algorithm. Energy-aware algorithms that depend on the first law of thermodynamics to minimize the overall DC energy budget [17-20] are effective in reducing cost and carbon footprint. An example is the *joint cooling and workload management* that simultaneously regulates the workload distribution and setpoint of the cooling equipment [21]. While the approach minimizes energy optimization, it does not address spatial overcooling since it does not account for the loss of available input cooling energy or irreversibility. Lowering

the irreversibility in a DC increases its exergetic efficiency and reduces overprovisioning of the cooling capacity.

The literature on irreversibility-based assessments of DCs is sparse. Component level exergy destruction has been considered for an open aisle room-based DC with perforated tiles [22-25]. The role of the number of active racks and hot and cold air mixing on energy and exergy efficiency has been identified [26], as well as the influence of hot air recirculation and cold air bypass [27, 28]. Computationally inexpensive reduced-order exergy models are available [29], for example, through flow network modeling (FNM) [30]. Although these exergy analyses consider the DC design, variations in operating parameters, and architecture-based thermal management [31], it is unclear how they can be used to explicitly assign server workload.

Therefore, we propose a novel approach for allocating workload that reduces both energy consumption and thermodynamic irreversibility inside air-cooled DCs. The proof of concept is demonstrated for a single rack DC with a rack-mountable cooling unit (RMCU). Workload distribution and cooling parameters are varied simultaneously, based on which a multi-objective genetic algorithm (MOGA) minimizes both energy and irreversibility and diminishes spatial overcooling. While existing workload scheduling algorithms use CFD and data-driven models, we employ an experimentally validated physics-based transient zonal model to characterize the flow and temperature field [32]. This improves the predictive extrapolation accuracy to determine a coarse-grained temperature distribution, which is the basis for simultaneously optimizing workload distribution and cooling.

Here, we (1) present a methodology to couple temperature predictions with a MOGA framework that optimizes energy and exergy efficiency, (2) examine the tradeoffs among energy consumption, irreversibility, cooling characteristics, and workload distribution, (3) compare decision variables from two single-objective optimization problems that either minimize energy consumption or irreversibility, (4) identify the benefits of solving a multi-objective optimization problem for different DC IT loads, (5) show the effect of overall IT load on the scheduling algorithm, and (6) generalize the results for the single rack DC architecture to multi-rack homogeneous DCs with RMCUs.

5.3 Methodology

5.3.1 Thermal model of the rack-based cooling infrastructure

We utilize a low complexity, spatiotemporal zonal model [32] to predict the temperature inside an enclosed single-rack micro DC with an RMCU [33]. The model provides spatial temperature predictions of an intermediate resolution, lying between a lumped-parameter thermal model and a complete CFD simulation. An enclosed rack with separated hot and cold aisles is considered, where the RMCU installed at the bottom of the rack supplies cold air to the IT equipment, as shown in Figure 5.1.

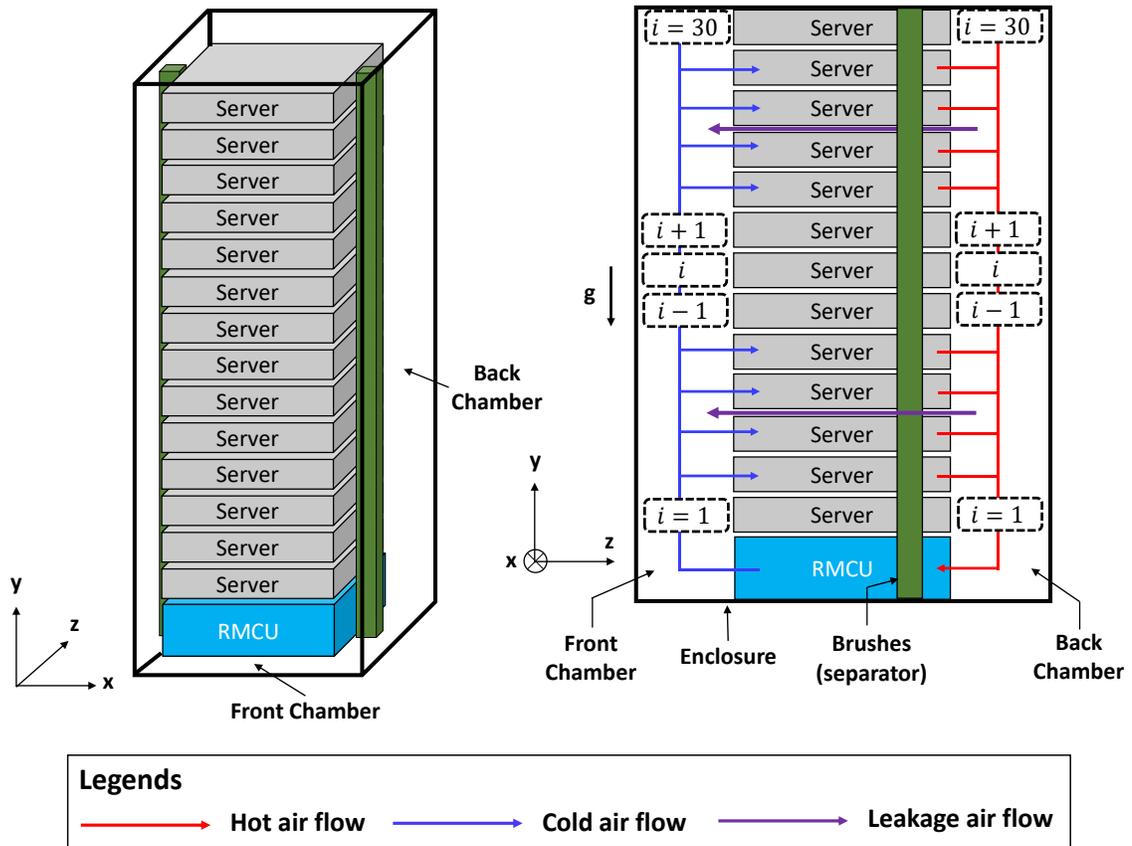


Figure 5.1: Schematic of the DC rack considered for the case study. The figure on the left depicts a three-dimensional view of the rack, while that on the right describes the salient airflows inside the enclosure. The IT rack consists of 30 1U (1U = 4.4 cm) servers and an RMCU.

The model predicts temperatures based on a zonal approach, where temperature, pressure, and flow rates are considered uniform within each zone. The zonal approach simplifies the spatial dependence of field variables and decreases the time required to execute the algorithm as compared to that required for full-field CFD simulations. Servers are considered to be variable heat sources with magnitudes that depend on their utilizations. Neglecting momentum and energy transfer to the ambient, six types of control volumes are identified, i.e., (1) zones lying in the fronts of servers, 2) zones at the backs of servers, 3) zones adjacent to the RMCU air supply, 4) zones adjacent to the RMCU air return, 5) each

server by itself, and 6) the RMCU. The cooling system, which consists of fans and a plate-fin air-water heat exchanger, utilizes chilled water from an external circulation loop. A vapor compression refrigeration (VCR) chiller with an ambient air-cooled condenser produces the required chilled water for the RMCU. Figure 5.2 is a schematic of the process flow. Temperature prediction in the architecture occurs through a two-step process, where an FNM based on the mechanical resistance provides the pressure and airflow in each zone inside the enclosure for use in the energy balance equations to obtain the zonal temperatures.

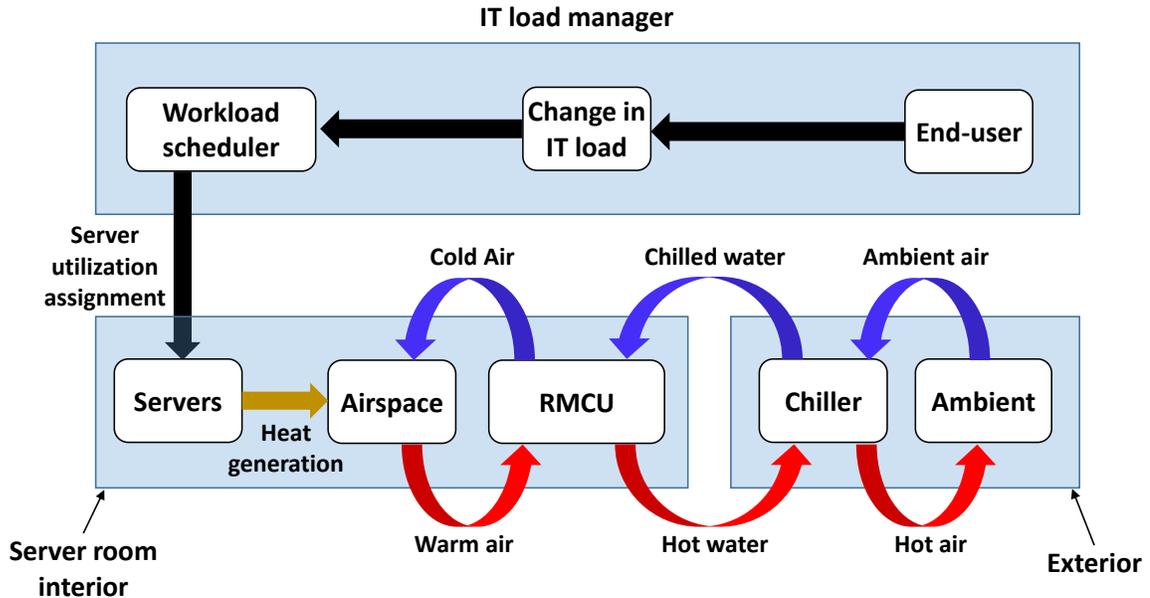


Figure 5.2: The process flow across different components for the DC case study. The IT rack shown in Figure 5.1 is situated within the server room interior block.

5.3.1.1 Step 1: Determination of the flowrates using mechanical resistances

Since the flow inside an enclosed DC is pressure-driven [32], the mechanical resistance circuit illustrated in Figure 5.3 is employed to calculate the flowrates and pressures. The RMCU is equivalent to a current source that is an airflow of magnitude \dot{Q}_a . Each active

server, either idle or utilized, is represented as a current source for an airflow \dot{Q}_i^s . Unpowered passive servers behave as porous separating media connecting the hot and cold chambers [32, 34] and are represented using a flow resistance R_s . The separators or brushes that prevent leakage airflow and energy transport across the chambers are represented by another resistance R_{br} . Since the principal direction of air travel inside the rack is the vertical direction, the flow resistance in the front chamber is denoted as $R_{f,v}$ and the corresponding resistance in the back chamber by $R_{b,v}$. The magnitudes of these resistances are obtained from our previous work [35]. The model assumes a large flow resistance across the IT rack and the ambient so that the air infiltration from the rack to ambient and in reverse is minimized [32]. This is readily ensured in practice by hindering leakages through the doors that enclose the IT rack by using sealing gaskets.

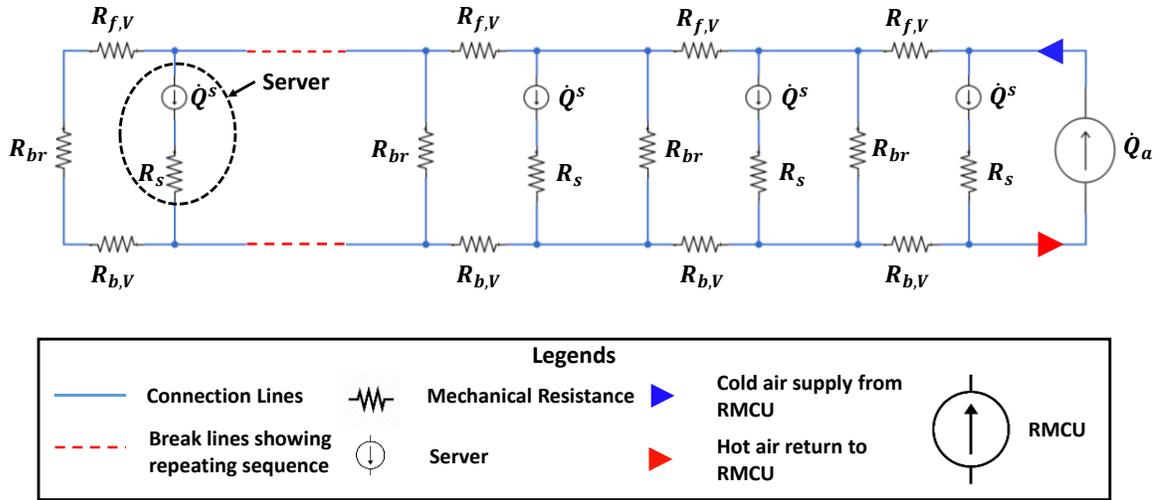


Figure 5.3: A flow resistance network representation for a server enclosure equipped with thirty servers and an RMCU. For simplification, the dotted red lines represent identical sequences of servers and mechanical resistances inside the enclosure.

The server flow rate is a function of the inlet air temperature of i^{th} server T_i^f [32],

$$\dot{Q}_i^s = \begin{cases} 0.01415 & \text{if } T_i^f < 25^\circ\text{C} \\ 0.01415 + (T_i^f - 25) \times 0.00142 & \text{if } 25^\circ\text{C} < T_i^f < 35^\circ\text{C} \end{cases} \quad (5.1)$$

Once the resistances and server flow rates are known, the mass balance equation for each server and its interacting zones can be constructed. Considering the mass balance in the front of an active server,

$$\left[\frac{P_{i+1}^f - P_i^f}{R_{f,V}} \right]^m + \left[\frac{P_{i-1}^f - P_i^f}{R_{f,V}} \right]^m + \left[\frac{P_i^b - P_i^f}{R_{br}} \right]^m - \dot{Q}_i^s = 0, \quad (5.2)$$

where m denotes the relationship between the flow rate and pressure, P pressure, i the vertical index of a server, f the front of a server, and b the back of that server.

The mass balance for the front zones of passive servers,

$$\left[\frac{P_{i+1}^f - P_i^f}{R_{f,V}} \right]^m + \left[\frac{P_{i-1}^f - P_i^f}{R_{f,V}} \right]^m + \left[\frac{P_i^b - P_i^f}{R_{br}} \right]^m + \left[\frac{P_i^b - P_i^f}{R_s} \right]^m = 0, \quad (5.3)$$

and for the zones at the backs of active servers,

$$\left[\frac{P_{i+1}^b - P_i^b}{R_{b,V}} \right]^m + \left[\frac{P_{i-1}^b - P_i^b}{R_{b,V}} \right]^m + \left[\frac{P_i^f - P_i^b}{R_{br}} \right]^m + \dot{Q}_i^s = 0. \quad (5.4)$$

The mass balance for the back zone of the passive server provides the relation,

$$\left[\frac{P_{i+1}^b - P_i^b}{R_{b,V}} \right]^m + \left[\frac{P_{i-1}^b - P_i^b}{R_{b,V}} \right]^m + \left[\frac{P_i^f - P_i^b}{R_{br}} \right]^m + \left[\frac{P_i^f - P_i^b}{R_s} \right]^m = 0, \quad (5.5)$$

and for zones in the front of the RMCU in the cold chamber,

$$\left[\frac{P_{i+1}^f - P_i^f}{R_{f,V}} \right]^m = \dot{Q}_a. \quad (5.6)$$

Utilizing the mass balance for zones in front of the RMCU in the hot chamber,

$$\left[\frac{P_{i+1}^b - P_i^b}{R_{b,V}} \right]^m = \dot{Q}_a. \quad (5.7)$$

Applying Eqs. (5.1)-(5.7) with $m = 1$ [35], leads to a linear relationship between the flow and pressure for each zone and results in a system of linear algebraic equations that provide the pressure corresponding to each zone.

5.3.1.2 Step 2: Zonal energy balance

Once the airflows and pressure for each zone are determined by solving the above set of linear equations, the energy balance formulations provide the zonal temperatures. The transient zonal energy balance equation for an active server is [32, 36],

$$\frac{X_s}{2} \left(\frac{dT_i^e}{dt} + \frac{dT_i^f}{dt} \right) = \rho_a C_{p,a} \dot{Q}_i^s (T_i^f - T_i^e) + \dot{P}_i^s, \quad (5.8)$$

where T_i^e denotes the exhaust temperature of the i^{th} server, T_i^f the front zone temperature of the i^{th} server, ρ_a the density of air, $C_{p,a}$ the specific heat of air, X_s the thermal mass of a server obtained from the literature [37], and \dot{P}_i^s the power consumption of the i^{th} server, which depends on its utilization [38-40].

An energy balance for the waterside within the RMCU results in [32],

$$\rho_w C_{p,w} V_w \left(\frac{dT_{c,w}}{dt} + \frac{dT_{h,w}}{dt} \right) = \rho_w \dot{Q}_w C_{p,w} (T_{c,w} - T_{h,w}) + \frac{UA}{2} (T_{h,a} + T_{c,a} - T_{c,w} - T_{h,w}), \quad (5.9)$$

and for the airside within the RMCU,

$$\rho_a C_{p,a} V_a \left(\frac{dT_{c,a}}{dt} + \frac{dT_{h,a}}{dt} \right) = \rho_a \dot{Q}_a C_{p,a} (T_{h,a} - T_{c,a}) - \frac{UA}{2} (T_{h,a} + T_{c,a} - T_{c,w} - T_{h,w}), \quad (5.10)$$

where $T_{c,a}$ and $T_{h,a}$ denote the cold side and hot side air temperatures of the heat exchanger inside the RMCU, \dot{Q}_a and \dot{Q}_w the air and water flow respectively, $T_{h,w}$ and $T_{c,w}$ the hot water outlet and chilled water inlet temperature, respectively, ρ_w the density of water, $C_{p,w}$ the specific heat of water, U the overall heat transfer coefficient of the heat exchanger as a

function of \dot{Q}_w and \dot{Q}_a [32], A the contact surface area responsible for heat transfer, and V_w and V_a the volumes of the water and air inside the heat exchanger, respectively.

Table 5.1: Expressions for φ_1 through φ_5 in Eq. (5.11)

Zones in front of the servers		Zones at the back of the servers	
Φ_1 (Flow transport in front chamber)		Φ_1 (Flow transport in back chamber)	
$[P_{i+1}^f - P_i^f] \geq 0$	$\rho_a C_{p,a} \left[\frac{P_{i+1}^f - P_i^f}{R_{f,v}} \right] T_{i+1}^f$	$[P_{i+1}^b - P_i^b] \geq 0$	$\rho_a C_{p,a} \left[\frac{P_{i+1}^b - P_i^b}{R_{b,v}} \right] T_{i+1}^b$
$[P_{i+1}^f - P_i^f] < 0$	$\rho_a C_{p,a} \left[\frac{P_{i+1}^f - P_i^f}{R_{f,v}} \right] T_i^f$	$[P_{i+1}^b - P_i^b] < 0$	$\rho_a C_{p,a} \left[\frac{P_{i+1}^b - P_i^b}{R_{b,v}} \right] T_i^b$
Φ_2 (Flow transport in front chamber)		Φ_2 (Flow transport in back chamber)	
$[P_{i-1}^f - P_i^f] \geq 0$	$\rho_a C_{p,a} \left[\frac{P_{i-1}^f - P_i^f}{R_{f,v}} \right] T_{i-1}^f$	$[P_{i-1}^b - P_i^b] \geq 0$	$\rho_a C_{p,a} \left[\frac{P_{i-1}^b - P_i^b}{R_{b,v}} \right] T_{i-1}^b$
$[P_{i-1}^f - P_i^f] < 0$	$\rho_a C_{p,a} \left[\frac{P_{i-1}^f - P_i^f}{R_{f,v}} \right] T_i^f$	$[P_{i-1}^b - P_i^b] < 0$	$\rho_a C_{p,a} \left[\frac{P_{i-1}^b - P_i^b}{R_{b,v}} \right] T_i^b$
Φ_3 (Flow transport through brushes in front chamber)		Φ_3 (Flow transport through brushes in back chamber)	
$[P_i^b - P_i^f] \geq 0$	$\rho_a C_{p,a} \left[\frac{P_i^b - P_i^f}{R_{br}} \right] T_i^b$	$[P_i^f - P_i^b] \geq 0$	$\rho_a C_{p,a} \left[\frac{P_i^f - P_i^b}{R_{br}} \right] T_i^f$
$[P_i^b - P_i^f] < 0$	$\rho_a C_{p,a} \left[\frac{P_i^b - P_i^f}{R_{br}} \right] T_i^f$	$[P_i^f - P_i^b] < 0$	$\rho_a C_{p,a} \left[\frac{P_i^f - P_i^b}{R_{br}} \right] T_i^b$
Φ_4 (Flow through passive servers)		Φ_4 (Flow through passive servers)	
$[P_i^b - P_i^f] \geq 0$	$\rho_a C_{p,a} \left[\frac{P_i^b - P_i^f}{R_s} \right] T_i^b$	$[P_i^f - P_i^b] \geq 0$	$\rho_a C_{p,a} \left[\frac{P_i^f - P_i^b}{R_s} \right] T_i^f$
$[P_i^b - P_i^f] < 0$	$\rho_a C_{p,a} \left[\frac{P_i^b - P_i^f}{R_s} \right] T_i^f$	$[P_i^f - P_i^b] < 0$	$\rho_a C_{p,a} \left[\frac{P_i^f - P_i^b}{R_s} \right] T_i^b$
Φ_5 (Flow through active servers)		Φ_5 (Flow through active servers)	
$-\rho_a C_{p,a} \dot{Q}_i^s T_i^f$		$\rho_a C_{p,a} \dot{Q}_i^s T_i^e$	

Applying the energy balance to the cold and hot chamber zones inside the enclosure,

$$\rho_a C_{p,a} V_z \alpha \left(\frac{dT_i}{dt} \right) = \varphi_1 + \varphi_2 + \varphi_3 + \varphi_4 + \varphi_5, \quad (5.11)$$

where α denotes the thermal mass correction factor, V_z the volume of a zone, and T_i the temperature of a zone. The expressions for φ_1 through φ_5 are reported in Table 5.1.

5.3.2 Energy consumption and power usage effectiveness

The server power consumption is a major contributor to the overall energy consumed by a DC. It consists of the energy consumed by the (1) CPUs and (2) server fans, where the second makes a small 0.95% contribution towards the overall power consumed by a server. Changing the server inlet temperature from 15 °C to 27 °C has been shown to increase the energy consumption of a server by only 0.48% [41]. Therefore, the power consumed by a server is solely a function of its utilization, i.e., the magnitude to which it is stressed [38-40]. The total IT power consumption,

$$\dot{P}_{IT} = \sum_{i=1}^N \dot{P}_i^s = \sum_{j=1}^N [C_1 + C_2 u_i^s], \quad (5.12)$$

where \dot{P}_{IT} represents the overall power consumption of the servers, u_i^s and \dot{P}_i^s the utilization level and power consumption of a server with the vertical index i , and C_1 and C_2 are constants that depend on the type and model of the server. We assume homogeneous server types within a DC rack so that C_1 and C_2 are constant throughout. For an HP ProLiant DL360 G5 server, their values are obtained from a datasheet $C_1 = 223.41$ and $C_2 = 154.5$ [42]. For an idle server, $u_i^s = 0$, $\dot{P}_i^s = C_1$ is the static power. For a passive server, $\dot{P}_i^s = 0$.

The overall power consumption in the cooling system,

$$\dot{P}_{cool} = \dot{P}_{ch} + \dot{P}_F, \quad (5.13)$$

where, \dot{P}_{ch} and \dot{P}_F denote the power consumed by the VCR chiller and the fans inside the RMCU, respectively. These two components of power consumption are calculated using Eqs. (5.16) and (5.17) below.

$$y = A_1x_1 + A_2x_2 + A_3x_3, \quad (5.14)$$

$$y = \left[\left(1 + \frac{1}{COP} \right) \frac{T_{c,w}}{T_{amb}} \right], \quad x_1 = \frac{T_{c,w}}{Q_{ch}}, \quad x_2 = \frac{T_{amb} - T_{c,w}}{Q_{ch}T_{amb}}, \quad x_3 = \frac{\left(1 + \frac{1}{COP} \right) Q_{ch}}{T_{amb}}. \quad (5.15)$$

$$\dot{P}_{ch} = \frac{[Q_{ch} + A_1T_{c,w}T_{amb} + A_2(T_{amb} - T_{c,w})]}{T_{c,w} - (A_3Q_{ch})} - Q_{ch}. \quad (5.16)$$

Eqs. (5.14) and (5.15) are from the Ng-Gordon universal chiller model [43], where $T_{c,w}$ represents the chilled water temperature leaving the evaporator, T_{amb} the ambient air temperature entering the condenser, Q_{ch} the heat load on the chiller, i.e., the IT load, and COP the coefficient of performance of the chiller based on the cooling load per unit power consumption. In this model, all temperatures are in K and Q_{ch} is in kW. The constants A_1 through A_3 are determined from the performance of a commercially available chiller [44], where multivariate linear regression results in $A_1 = 0.0092$, $A_2 = 8.953$, and $A_3 = 0.649$. The value of T_{amb} is set to 297.039 K as per the chiller performance data. Eq. (5.16) describes the power consumption of the chiller by combining Eqs. (5.14) and (5.15).

The total power consumption of the fans inside the RMCU,

$$\dot{P}_F = N_F \left[B_1 + B_2 \left(\frac{\dot{Q}_a}{N_F} \right) \right], \quad (5.17)$$

where N_F denotes the number of fans and \dot{Q}_a the total airflow through RMCU. For our case, $N_F = 5$. The constants B_1 and B_2 are determined through experiments for the counter-rotating axial fan San Ace 80 9CRB [45], resulting in $B_1 = -0.268$ and $B_2 = 4.12 \times$

10^{-3} . While fan models can be based on linear, parabolic, and cubic relations between the flowrate and power consumption, the linear relationship between airflow and energy consumption holds for our operational range of $350 \leq \dot{Q}_a \leq 850$ CFM.

The IT and cooling power consumption are combined together using the widely accepted energy-based metric, power usage effectiveness *PUE* [46-49],

$$PUE = 1 + \frac{\dot{P}_{cool}}{\dot{P}_{IT}}, \quad (5.18)$$

where $1 \leq PUE < \infty$. The *PUE* must be minimized, where the ideal but unrealistic value $PUE = 1$ implies that the DC consumes none of the total energy consumption for cooling purposes.

5.3.3 Exergy destruction and the 2nd law efficiency

The components of overall exergy destruction are the exergy losses (1) $\dot{\psi}_F$ in fans due to work transfer, (2) $\dot{\psi}_{hx}$ due to heat transfer inside the RMCU, (3) $\dot{\psi}_{ch}$ in the chiller, and (4) $\dot{\psi}_s$ in the servers due to heat dissipation. Here, $\dot{\psi}_s$ depends on the IT load imposed on the rack and end-user computing demand, which the DC administrator has no control over. Hence, overall exergy destruction in the optimization framework considers only the first three components that represent cooling cycle inefficiency [31],

$$\dot{\psi}_{cool} = \dot{\psi}_F + \dot{\psi}_{ch} + \dot{\psi}_{hx}. \quad (5.19)$$

The exergy loss associated with the fan is solely a function of airflow [30], i.e.,

$$\dot{\psi}_F = (1 - \eta_F)N_F \left[B_1 + B_2 \left(\frac{\dot{Q}_a}{N_F} \right) \right], \quad (5.20)$$

where the efficiency of the fan, $\eta_F = 0.7$, is obtained from a datasheet [45].

The exergy loss due to heat transfer inside the RMCU is determined using the relation [31],

$$\begin{aligned} \dot{\psi}_{hx} = & \rho_a \dot{Q}_a C_{p,a} \left[(T_{h,a} - T_{c,a}) - T_{amb} \ln \left(\frac{T_{h,a}}{T_{c,a}} \right) \right] \\ & + \rho_w \dot{Q}_w C_{p,w} \left[(T_{c,w} - T_{h,w}) - T_{amb} \ln \left(\frac{T_{c,w}}{T_{h,w}} \right) \right], \end{aligned} \quad (5.21)$$

where $T_{h,a}$ and $T_{c,a}$ denote the average hot air return temperature and cold air supply temperature of the RMCU, $T_{c,w}$ and $T_{h,w}$ are the chilled water supply and hot water return temperatures to the RMCU.

The exergy destruction in the chiller [31],

$$\dot{\psi}_{ch} = \rho_w \dot{Q}_w C_{p,w} T_{amb} \ln \left(\frac{T_{c,w}}{T_{h,w}} \right) + \rho_a \dot{Q}_{cf} C_{p,a} T_{amb} \ln \left(\frac{T_{h,amb}}{T_{amb}} \right), \quad (5.22)$$

where $T_{h,amb}$ denotes the warm air temperature at the exhaust of the VCR chiller condenser, T_{amb} the ambient air temperature supplied to the condenser, \dot{Q}_w the total volume flow rate through the chilled water loop, and \dot{Q}_{cf} the airflow of the condenser fans. The value of \dot{Q}_{cf} is obtained from the datasheet for the chiller [44].

The second law efficiency η_{2nd} is the exergy efficiency, which is the extent of irreversibility in the system as a fraction of total input exergy, i.e.,

$$\eta_{2nd} = \left(1 - \frac{\dot{\psi}_{cool}}{\dot{\psi}_{in}} \right) \times 100. \quad (5.23)$$

The exergy input $\dot{\psi}_{in}$ to the cooling system [31],

$$\dot{\psi}_{in} = \dot{P}_{ch} + \dot{P}_F + \rho_a \dot{Q}_a C_{p,a} \left[(T_{h,a} - T_{amb}) - T_{amb} \ln \left(\frac{T_{h,a}}{T_{amb}} \right) \right]. \quad (5.24)$$

From a thermodynamic perspective, system irreversibilities must be reduced so that as much of the available input energy can be used as is possible. This can be monitored by comparing η_{2nd} for different scenarios.

5.3.4 Optimization problem formulation

Using the models in Sections 5.3.2 and 5.3.3, our objectives are to (1) minimize PUE and (2) maximize η_{2nd} for the specific DC configuration considered. To examine the interplay between these two objectives, we consider three different optimization problems that are described in Table 5.2, where (1) the first problem minimizes PUE , (2) the second maximizes η_{2nd} , and (3) the third simultaneously minimizes PUE and maximizes η_{2nd} . The primary deciding factors that regulate PUE and η_{2nd} are (1) $T_{c,w}$, (2) \dot{Q}_a and (3) the server utilizations u_i^s values for which are optimized. Two classes of servers are considered here, (1) active servers with $0 < u_i^s \leq 1$ and (2) idle servers with $u_i^s = 0$. Turning OFF an idle server, i.e., making it a passive server to save energy, changes the thermal and optimization framework significantly [21, 32, 35]. We maintain the maximum server inlet temperature below the ASHRAE guideline temperature $T_g = 27^\circ\text{C}$ [50].

Table 5.2: Mathematical structure of three different optimization problems

Optimization 1	Optimization 2	Optimization 3
$\underset{u_i^s, T_{c,w}, Q_a}{\text{minimize}} \quad PUE = 1 + \frac{\dot{P}_{cool}}{\dot{P}_{IT}}$	$\underset{u_i^s, T_{c,w}, Q_a}{\text{maximize}} \quad \eta_{2nd} = \left(1 - \frac{\dot{\psi}_{cool}}{\dot{\psi}_{in}} \right)$	$\underset{u_i^s, T_{c,w}, Q_a}{\text{minimize}} \quad PUE$ $\underset{u_i^s, T_{c,w}, Q_{cu}}{\text{maximize}} \quad \eta_{2nd}$
$\text{s.t.} \quad \sum_{i=1}^n u_i^s = D$ $\max(T_i^f) \leq T_g$ $10 \leq T_{c,w}(\text{°C}) \leq 22^\circ\text{C}$ $350 \leq \dot{Q}_a(\text{CFM}) \leq 850$ $0 \leq u_i^s \leq 1, i = 1, \dots, n$	$\text{s.t.} \quad \sum_{i=1}^n u_i^s = D$ $\max(T_i^f) \leq T_g$ $10 \leq T_{c,w}(\text{°C}) \leq 22^\circ\text{C}$ $350 \leq \dot{Q}_a(\text{CFM}) \leq 850$ $0 \leq u_i^s \leq 1, i = 1, \dots, n$	$\text{s.t.} \quad \sum_{i=1}^n u_i^s = D$ $\max(T_i^f) \leq T_g$ $10 \leq T_{c,w}(\text{°C}) \leq 22^\circ\text{C}$ $350 \leq \dot{Q}_a(\text{CFM}) \leq 850$ $0 \leq u_i^s \leq 1, i = 1, \dots, n$

Optimization problems 1 and 2 are single objectives and multi-dimensional with both linear and nonlinear constraints, whereas the third optimization problem is multi-objective. The

primary constraint is a linear constraint that signifies that the sum of all u_i^s in the DC should equal the total offered workload D . The secondary constraint is non-linear, which maintains the ASHRAE thermal reliability guideline temperature of 27°C within the cold chamber. Evaluating the maximum temperature in the cold aisle requires the use of the coupled FNM and thermal model. The first two problems are solved using the MATLAB *fmincon* function with a *sequential quadratic programming* solver, whereas a genetic algorithm-based multi-objective solver *gamultiobj* is used to solve the third problem. Both functions support linear and nonlinear constraints. The lower and upper bounds for each decision variable (see Table 5.2) are obtained from the component datasheets of the RMCU [33] and the VCR chiller [44]. The convergence criterion for the objective functions is set to 10^{-6} for all three optimization problems.

5.4 Results and discussion

5.4.1 Temperature nonuniformity in the cold aisle

In an air-cooled DC, altering the cooling parameters such as $T_{c,w}$, \dot{Q}_a , and \dot{Q}_w significantly change the (1) flow-field, (2) temperature distribution, (3) energy consumption, and (4) system irreversibility, which leads to tradeoffs across different decision variables. Therefore, we examine the influence of these parameters on the cold aisle temperature distribution. Table 5.3 lists different cases for a specific workload distribution inside the contained single-rack DC. The two primary decision variables $T_{c,w}$ and \dot{Q}_a influence the energy consumptions of the chiller and fans, respectively. The effect of altering the water flow is neglected because the pumping power is only 2-3% of the overall cooling power

[8, 47, 51, 52]. For instance, changing the water flow by ~10% alters the air temperature by 3% [32, 35]. The model considers a fixed speed pump with an adequate valve mechanism-based flowrate control that delivers $\dot{Q}_w = 20$ GPM. As an alternative, a variable speed pump with a variable frequency drive can be used to control the energy consumption of the pump by regulating \dot{Q}_w .

Table 5.3: Case study parameters for the results presented in Figure 5.4, where variables $T_{c,w}$ and \dot{Q}_a are changed to constrain the maximum temperature in the cold aisle below 27 °C.

Case	$T_{c,w}$ (°C)	\dot{Q}_a (CFM)	Workload assignment with $D = 15$	Maximum allowable temperature in the cold chamber (°C)	Variance in temperature distribution	Water flowrate (GPM)
1	10	578	Each server has 50% utilization	27	11.6	20
2	12	609			8.6	
3	14	643			6.1	
4	16	682			3.9	
5	18	726			2.2	
6	20	776			1.0	
7	22	834			0.2	

Figure 5.4 presents the cold aisle temperature distributions corresponding to the cases listed in Table 5.3, where $T_{c,w}$ and \dot{Q}_a are varied to constrain the maximum temperature in the cold aisle below 27 °C. For a fixed workload distribution with all servers at 50% utilization, increasing both $T_{c,w}$ and \dot{Q}_a , increases the temperature uniformity in the cold aisle as reflected by the temperature variance reported in Table 5.3.

The relative magnitudes of natural and forced convection in the DC are characterized through the Richardson number $Ri = g\beta_a(T_{max}^f - T_{min}^f)H/V_c^2$ [53]. Natural convection is negligible compared to forced convection when $Ri \leq 0.1$. For the DC, $0.004 \leq Ri \leq 0.07$, implying that buoyancy can be neglected for the pressure-driven flow with forced convection.

The maximum air temperature in the cold chamber is observed in front of the server located furthest from the RMCU. This observation is consistent with the FNM depicted in Figure 5.3. The flow resistance in the front and back chambers at the i^{th} zone in the chamber is a function of its distance from the RMCU. An increase in the airflow path produces a higher flow resistance and increases the pressure drop [8], which decreases the transport of cold air that proceeds towards the top of the enclosed rack. Therefore, the server located furthest from the RMCU now has the highest intake air temperature. Comparison of cases 1 through 7, shows that the extent of overcooling decreases when $T_{c,w}$ and \dot{Q}_a are simultaneously increased.

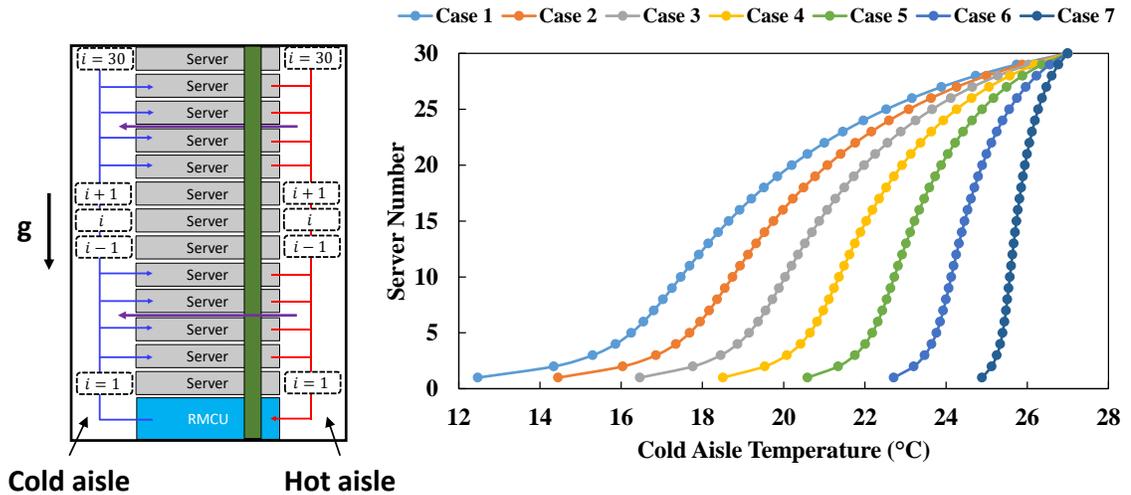


Figure 5.4: Nonuniformity in the cold aisle temperature due to changes in $T_{c,w}$ and \dot{Q}_a for a fixed workload distribution with all servers at 50% utilization for the cases listed in Table 5.3.

The impact of redistributing IT load is also investigated, where Table 5.4 lists three cases that are considered. The overall offered workload, D is set to 15, which can be achieved through numerous possible combinations of server utilizations. A fixed speed pump delivers $\dot{Q}_w = 20$ GPM and with $T_{c,w}$ fixed at 16 °C.

Table 5.4: Case study parameters for the results presented in Figure 5.5, where \dot{Q}_a and workload distributions are changed to constrain the maximum temperature in the cold aisle below 27 °C.

Case	$T_{c,w}$ (°C)	\dot{Q}_a (CFM)	Workload distribution with $D = 15$	Maximum allowable temperature in the cold chamber (°C)	Variance in temperature distribution	Water flow rate (GPM)
1	16	605	Workload assigned towards RMCU and $(u_{i=1 \rightarrow 15}^s = 1$ and $u_{i=16 \rightarrow 30}^s = 0)$	27	3.0	20
2	16	682	Uniform workload assignment $(u_{i=1 \rightarrow 30}^s = 0.5)$		3.9	
3	16	727	Workload assigned far from RMCU $(u_{i=1 \rightarrow 15}^s = 0$ and $u_{i=16 \rightarrow 30}^s = 1)$		4.4	

For the three IT load distributions, \dot{Q}_a is regulated to always maintain the maximum temperature in the cold aisle below 27 °C to ensure safe IT equipment operation. Figure 5.5

presents cold temperature distributions for three workload assignments. Server utilization for these cases are (1) $u_{i=1 \rightarrow 15}^s = 1$ and $u_{i=16 \rightarrow 30}^s = 0$, (2) $u_{i=1 \rightarrow 30}^s = 0.5$, and (3) $u_{i=1 \rightarrow 15}^s = 0$ and $u_{i=16 \rightarrow 30}^s = 1$. In practice, the workload distribution (or server utilization) is controlled through workload manager software. Assigning the IT workload towards the RMCU (Case 1) results in a more uniform temperature distribution for even lower \dot{Q}_a . Displacing the workload to servers that lie further away from the RMCU (Case 3) results in a higher variance (Table 5.4) in the temperature distribution across the rack and increases \dot{Q}_a . Consequently, the power consumption of the fans for Case 3 also increases. Therefore, displacing the workload towards the RMCU offers the potential to reduce overcooling and lower fan power consumption.

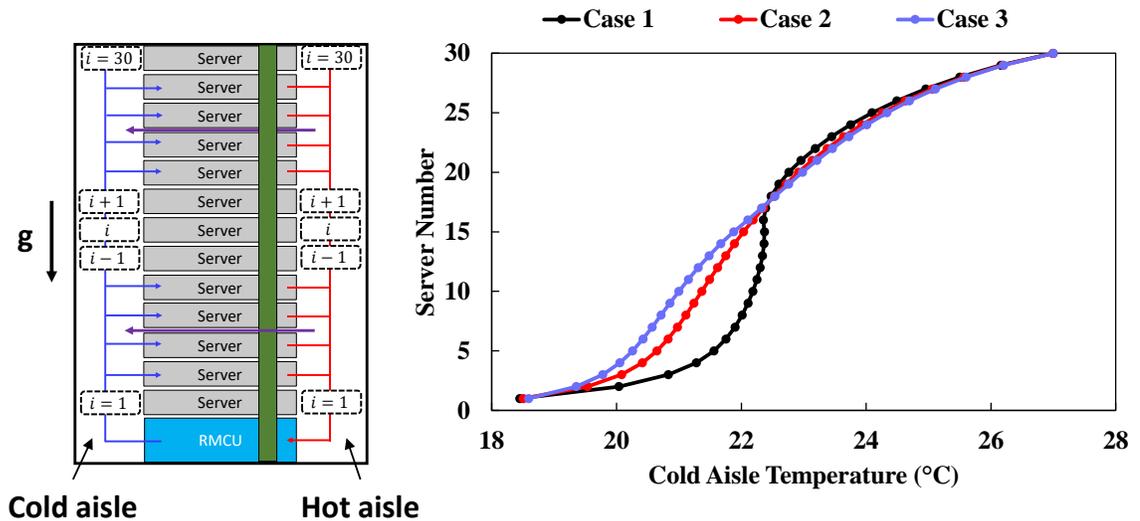


Figure 5.5: Non-uniformity in the cold aisle temperature due to simultaneous change in cooling parameter and workload distribution. The figure is to be interpreted in accordance with Table 5.4.

5.4.2 Impact on PUE and η_{2nd} of changing $T_{c,w}$, \dot{Q}_a , and workload distribution

To investigate the extent of overcooling when $T_{c,w}$, \dot{Q}_a , and the IT load distributions are changed, we consider two dimensionless metrics. The first PUE , defined in Eq. (18), is the overall power consumption as a fraction of the IT equipment power consumption. The exergy efficiency of the cooling system, based on the second law of thermodynamics, is expressed by Eq. (23) and accounts for the loss of available cooling energy due to inherent system irreversibilities.

We define a non-dimensional number $\theta \times \beta$ where,

$$\theta = \frac{T_{c,w}}{T_{amb}}, \quad \beta = \frac{\dot{Q}_a}{\dot{Q}_{IT}} = \frac{\dot{Q}_a}{125 \times \dot{P}_{IT}}. \quad (5.25)$$

Here, θ is the ratio of chilled water temperature to the ambient dead state temperature, and β compares the RMCU airflow \dot{Q}_a with the total airflow requirement for all servers \dot{Q}_{IT} . A large value of θ lowers the chiller power consumption and decreases heat transfer irreversibility between the cold water and the ambient dead state. Although the total airflow through the servers as a function of inlet air temperature is provided by Eq. (1), we utilize a standard DC guideline that prescribes 125 CFM airflow to be provided for each kW IT load increment [54], i.e., $\dot{Q}_{IT} = 125 \times \dot{P}_{IT}$, to ensure the thermal reliability of servers. A larger β results in higher fan power consumption and lowers system irreversibility by diminishing the temperature gradients in the airspace. The impact of changing $\theta \times \beta$ on the PUE and η_{2nd} for different maximum allowable air temperatures in the cold aisle, i.e., 26, 27, and 28 °C is illustrated in Figure 5.6. For all cases, the total workload assignment $D = \sum_{i=1}^{30} u_i$ is set to 15, where each of the 30 servers inside the rack is 50% utilized.

Examination of Figure 5.6 shows that increasing $\theta \times \beta$ increases PUE and also improves η_{2nd} . If $T_{c,w}$ is increased, \dot{Q}_a must also be simultaneously increased to satisfy the maximum allowable cold aisle temperature T_{cr} . As a result, the chiller power consumption decreases, and that of the fans increases. Since fan operation is costlier than chiller operation for modular DCs [44, 45], increasing $\theta \times \beta$ increases PUE . On the other hand, from a second law perspective, increasing θ reduces the heat transfer irreversibility in the chilled water stream, and higher β reduces the irreversibility due to spatial air temperature gradients. Therefore, with an increase in $\theta \times \beta$ the η_{2nd} improves. Hence, larger values of η_{2nd} are desirable for minimizing the system irreversibility, but this increases PUE , increasing the cooling power consumption, which is undesirable.

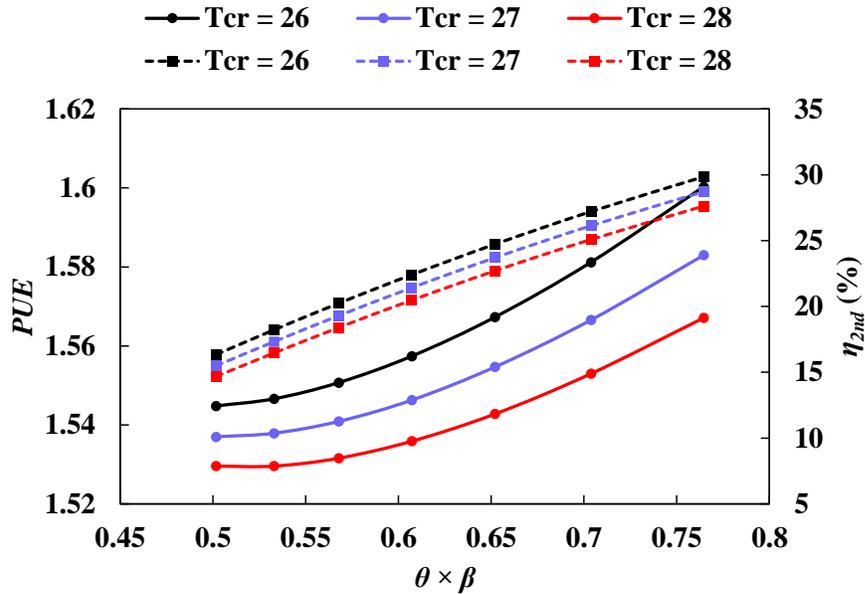


Figure 5.6: Effect of cooling system operating parameters on PUE and η_{2nd} for different maximum allowable temperatures in the cold aisle. The solid lines with circles represent variations in PUE , whereas the dashed lines with square symbols represent changing η_{2nd} .

Figure 5.7 presents the tradeoffs in PUE and η_{2nd} as a function of $\theta \times \beta$ for different workload distributions when the maximum air temperature in the cold aisle is maintained below 27 °C. The total offered workload $D = \sum_{i=1}^{30} u_i = 15$. Increasing $\theta \times \beta$ increases both η_{2nd} and PUE . This can be addressed by changing the workload distribution, where three cases are considered, i.e., (1) the workload is displaced towards the RMCU ($u_{i=1 \rightarrow 15}^s = 1$ and $u_{i=16 \rightarrow 30}^s = 0$), (2) there is uniform workload distribution with 50% utilization ($u_{i=1 \rightarrow 30}^s = 0.5$), and (3) the workload is displaced away from the RMCU ($u_{i=1 \rightarrow 15}^s = 0$ and $u_{i=16 \rightarrow 30}^s = 1$). The lowest PUE is obtained for the first case when the workload is displaced towards the RMCU, but this also leads to the lowest η_{2nd} since system irreversibility increases. When the workload is displaced away from the RMCU for Case 3, PUE increases by ~5% as compared to Case 1 since the flow path resistance between the heat source and cooling unit increases. This increases the pressure drop, which resists the incoming cold air from the RMCU, a resistance that increases as $\theta \times \beta$ is increased. As a result, Case 3 alters the temperature distribution inside the DC slightly, leading to a small increase in PUE . Increasing $\theta \times \beta$ for this case reduces irreversibility within the DC airspace and chilled water stream, improving η_{2nd} . The tradeoff between increasing PUE and decreasing η_{2nd} implies that minimization of power consumption and irreversibility cannot be achieved simultaneously for all operational cases. Therefore, we solve three nonlinear optimization problems in the subsequent sections that (1) minimize PUE , (2) maximize η_{2nd} , and (3) simultaneously minimize PUE and maximize η_{2nd} through a tradeoff.

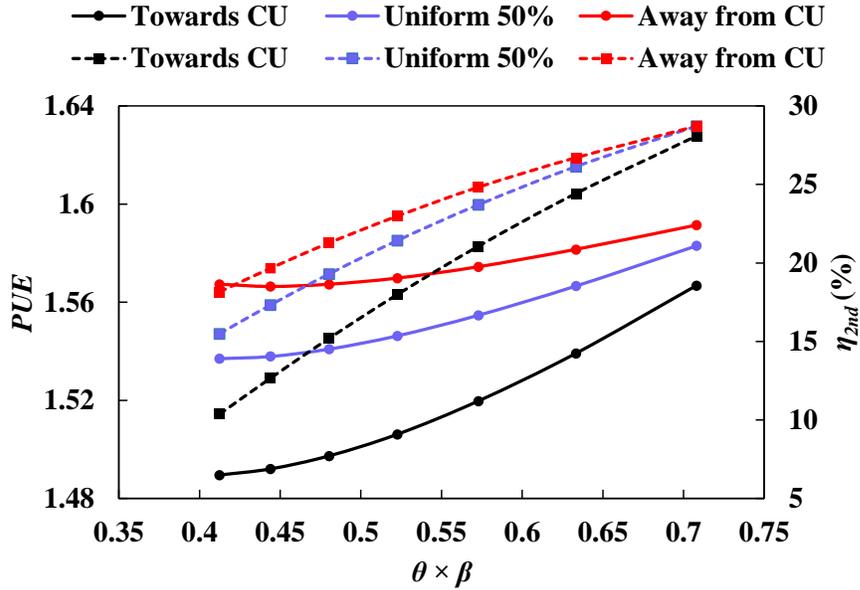


Figure 5.7: Effect of cooling system operation parameters on PUE and η_{2nd} for different workload distributions. The solid lines with circles represent variations of PUE whereas dashed lines with square symbols represent changes in η_{2nd} .

5.4.3 Energy optimization by minimizing PUE

The first optimization problem from Table 5.2 is the PUE minimization problem, which intrinsically minimizes cooling energy consumption. A thermal reliability constraint is imposed to limit the maximum temperature in the cold chamber to below 27 °C. Table 5.5 shows results for different D following the methodology presented in Section 2.4. Optimal values of PUE and the cooling decision variables $T_{c,w}$ and \dot{Q}_a are also presented in Table 5.5.

Table 5.5: Optimal solutions for the different *PUE* minimization cases presented in Figure 5.8.

Case	Total offered workload (D)	Airflow of RMCU \dot{Q}_a (CFM)	Chilled water temperature $T_{c,w}$ (°C)	Minimized <i>PUE</i>	Variance in cold aisle temperature distribution	Corresponding value of η_{2nd}
1	6.0	498	10	1.57	11.5	11.1
2	10.5	493	10	1.52	10.0	10.8
3	15.0	488	10	1.49	8.1	10.4
4	19.5	483	10	1.46	7.0	10.0
5	24.0	550	13.7	1.44	5.0	14.3

Figure 5.8 represents the optimal workload assignments and resulting cold aisle temperature distributions across the rack. As D increases, the additional servers that lie towards the top of the rack become utilized. The maximum achievable computing load $D_{max} = 30$ for this case. For the smallest value investigated, $D = 6$, servers adjacent to the RMCU are utilized since allocating workload closest to the cooling unit results in the lowest flow path resistance. This improves advective transport of the cold air in the front chamber, reducing overcooling of air by simultaneously regulating $T_{c,w}$ and \dot{Q}_a . Consequently, *PUE* is lowered across a wide range of D values. Since fan operation in the RMCU is more expensive than operating the chiller, low values of $T_{c,w}$ and \dot{Q}_a are favored to minimize overall energy consumption. An increase in D improves *PUE* since the load factor, $LF = D/D_{max}$ increases. As LF increases, the coefficient of performance (COP) of the chiller reaches its designed nominal capacity, and the energy efficiency of the cooling system improves. A higher LF also increases the uniformity in the cold aisle, increasing the mean air temperature. This increase in the average temperature results in a lower value

for $T_{amb} - T_{mean}^f$, reducing the system irreversibility due to heat transfer. Therefore, for the highest D , η_{2nd} improves by $\sim 4\%$, as shown in Table 5.5. Since the goal is to minimize PUE , the solution does not optimize exergy efficiency. The resulting tradeoffs in PUE and η_{2nd} are illustrated in Figures 5.6 and 5.7.

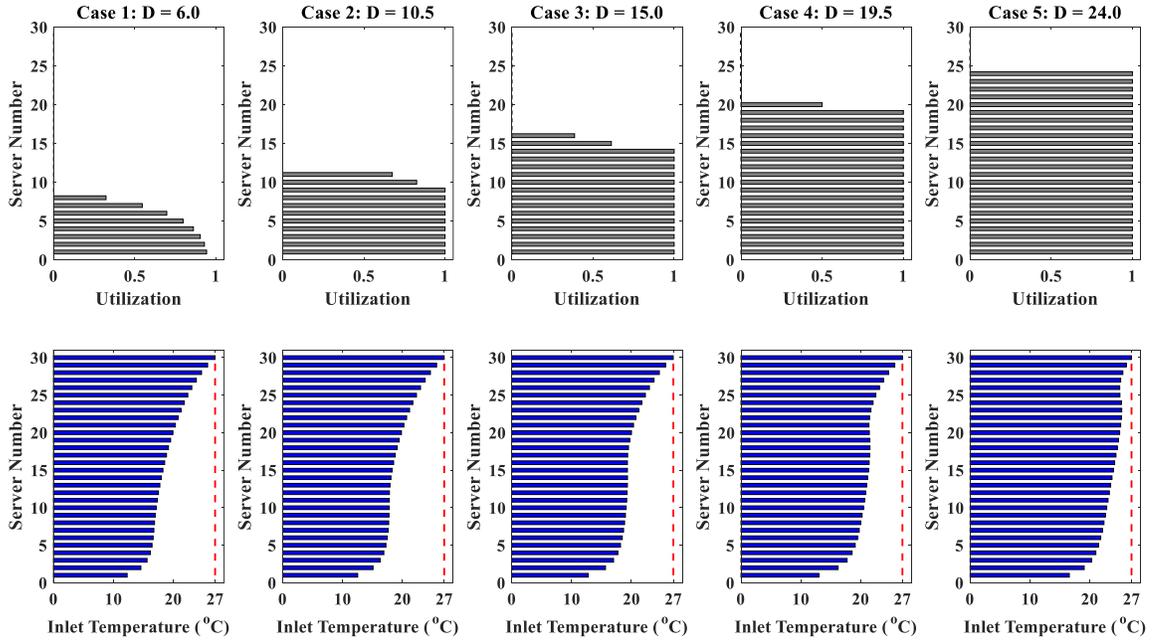


Figure 5.8: PUE minimized workload assignment and resultant cold aisle temperature distributions across the rack for different offered workloads D . The red dotted line corresponds to the maximum allowable server inlet air temperature, 27°C . The figure should be interpreted in accordance with Table 5.5.

5.4.4 Exergy optimization by maximizing η_{2nd}

The second optimization problem (see Table 5.2) is the η_{2nd} maximization problem, which reduces the system irreversibility. A nonlinear thermal reliability constraint is imposed that limits the maximum temperature in the cold chamber to below 27°C . The optimization problem is solved for the different D listed in Table 5.6 using the methodology in Section 2.4. The optimal values of η_{2nd} and cooling decision variables $T_{c,w}$ and \dot{Q}_a , are also

provided in Table 5.6. Figure 5.9 represents the optimal workload assignment obtained by solving the maximization problem and resulting cold aisle temperature distribution across the rack.

Table 5.6: Optimized solutions for different η_{2nd} maximization cases presented in Figure 5.9.

Case	Total offered workload D	\dot{Q}_a (CFM)	$T_{c,w}$ (°C)	Maximized η_{2nd}	Variance in cold aisle temperature distribution	Corresponding value of PUE
1	6.0	850	22	30.05	0.066	1.68
2	10.5	850	22	29.92	0.057	1.63
3	15.0	850	22	29.75	0.048	1.59
4	19.5	850	22	29.50	0.043	1.55
5	24.0	850	22	29.13	0.052	1.52

Similar to the PUE minimization problem, η_{2nd} maximization also assigns workload closest to the cooling unit for the wide range of D values. However, optimal values of the cooling decision variables $T_{c,w}$ and \dot{Q}_a are different from those obtained from the PUE minimization problem. Maximizing η_{2nd} increases $T_{c,w}$ and \dot{Q}_a , which reduces the (1) exergy destruction for the chilled water stream and (2) results in a uniform temperature distribution, as shown in Figure 5.9. Therefore, two major heat transfer irreversibilities in the system, i.e., $T_{amb} - T_{mean}^f$ and $T_{amb} - T_{c,w}$ are minimized. By minimizing the $T_{amb} - T_{mean}^f$, the irreversibility due to thermal gradients in the airspace is lowered. On the other hand, minimizing $T_{amb} - T_{c,w}$ reduces the irreversibility associated with the chilled water stream.

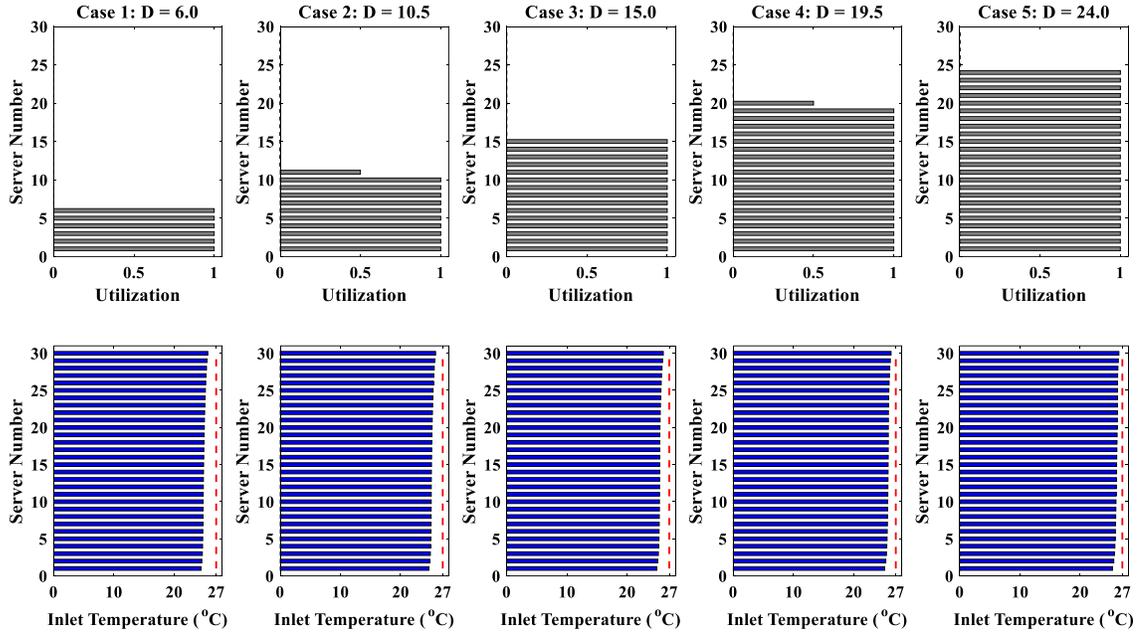


Figure 5.9: η_{2nd} maximized workload assignment and resulting cold aisle temperature distributions across the rack for different offered workloads D . The red dotted line corresponds to the maximum allowable server inlet air temperature, 27 °C. The figure should be interpreted in accordance with Table 5.6.

5.4.5 Multi-objective optimization

Figure 5.10 presents results from two different single-objective optimization problems as a function of LF for the DC. By minimizing PUE , the exergy efficiency of the system decreases by approximately 19%, which leads to a significant loss of available input cooling energy. To achieve maximum exergy efficiency, PUE increases by 7%, leading to higher energy consumption. The value of η_{2nd} for the PUE minimization problem increases by $\approx 4\%$ for $LF = 0.8$, as shown by the black dashed line in Figure 5.10. This improvement is attributed to the irreversibility decrease in the chilled water stream and a smaller spatial air temperature gradient shown in Table 5.5.

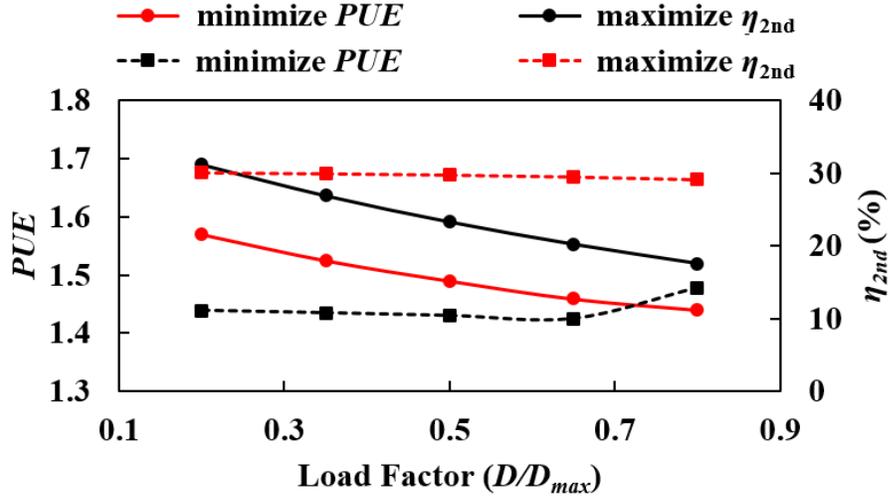


Figure 5.10: PUE and η_{2nd} as a function of LF for PUE minimization and η_{2nd} maximization problems. Solid lines with circles represent variations of PUE , whereas dashed lines with square symbols represent those for η_{2nd} . The optimization details for the figure are provided in Sections 5.3.3 and 5.3.4.

To address the tradeoff between PUE and η_{2nd} , a multi-objective optimization problem subject to the same set of constraints is considered. The structure of the third optimization problem is provided in Table 5.2. The solution is obtained using a genetic algorithm-based multi-objective solver where PUE is minimized and η_{2nd} is maximized simultaneously. The multi-objective problem is solved for the LF values specified in Tables 5.5 and 5.6. Figure 5.11 shows the Pareto optimal front of PUE vs. η_{2nd} obtained by solving the multi-objective optimization problem for five different LF values, 0.2, 0.35, 0.5, 0.65, and 0.8. Generally, the trends suggest that for all LF η_{2nd} improves at the cost of PUE . For a higher LF (marked as the desired region in Figure 5.11), PUE decreases while η_{2nd} improves so that increasing the network traffic load is a promising method to reduce the overcooling of a DC rack.

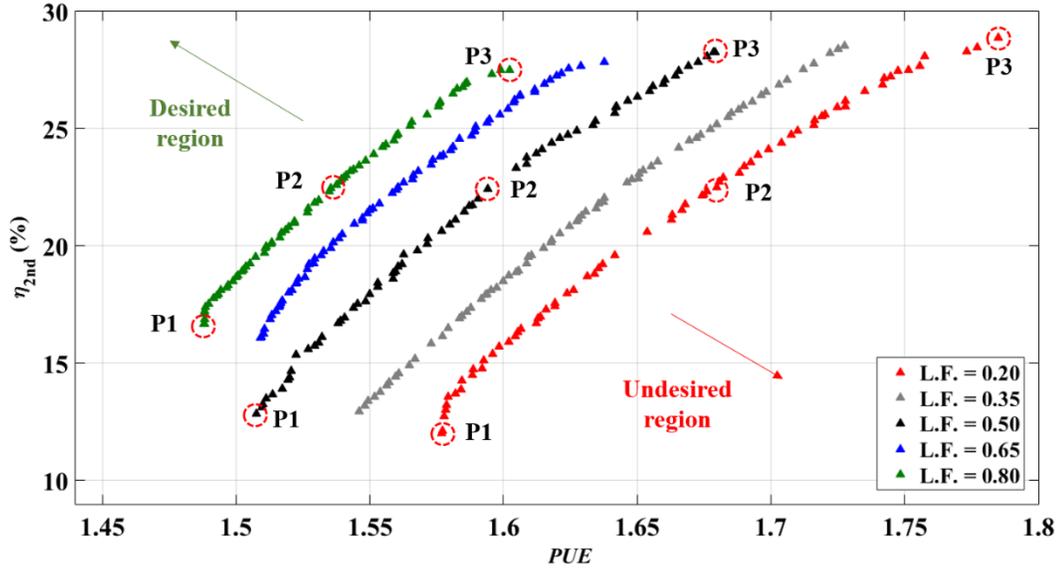


Figure 5.11: Pareto front of PUE and η_{2nd} obtained by solving the multi-objective optimization problem for five different LF s imposed on the DC. The symbols in the figure should be interpreted in accordance with Table 5.7 and Figure 5.12.

To gain better insight into the cooling characteristics and workload distribution in the Pareto front, nine points are marked (using red dotted circles) on the LF curves of 0.2, 0.5, and 0.8. in Figure 5.11. The corresponding cooling parameters, workload assignments, and temperature distributions are shown in Figure 5.12 and Table 5.7. For different LF s, optimizations 1 and 2 results in a utilization distribution that places the workloads closest to the RMCU (Figures 5.8 and 5.9). In contrast, the workload distribution obtained from optimization 3 (Figure 5.12) distributes the IT load across the rack in a nonuniform manner. Figure 5.8 shows that Optimization 1 overcools several servers in the vicinity of the RMCU. Therefore, distributing servers with high and low utilizations across the rack and thus regulating the cooling parameters should result in the desired tradeoff across the two objectives. By comparing three points for the same LF from Figure 5.12, we infer that moving across the Pareto front for a specific LF does not influence the spatial workload

distribution. The tradeoff in the Pareto front is a result of the synchronized regulation of \dot{Q}_a and $T_{c,w}$, which alters the extent of reduction in overcooling and changes the variance in the cold aisle temperature distribution (Table 5.7). Moving from left to right on the Pareto front for each LF curve results in a lower variance in temperature, which consequently improves η_{2nd} . However, such an improvement leads to additional expenditure from a PUE perspective.

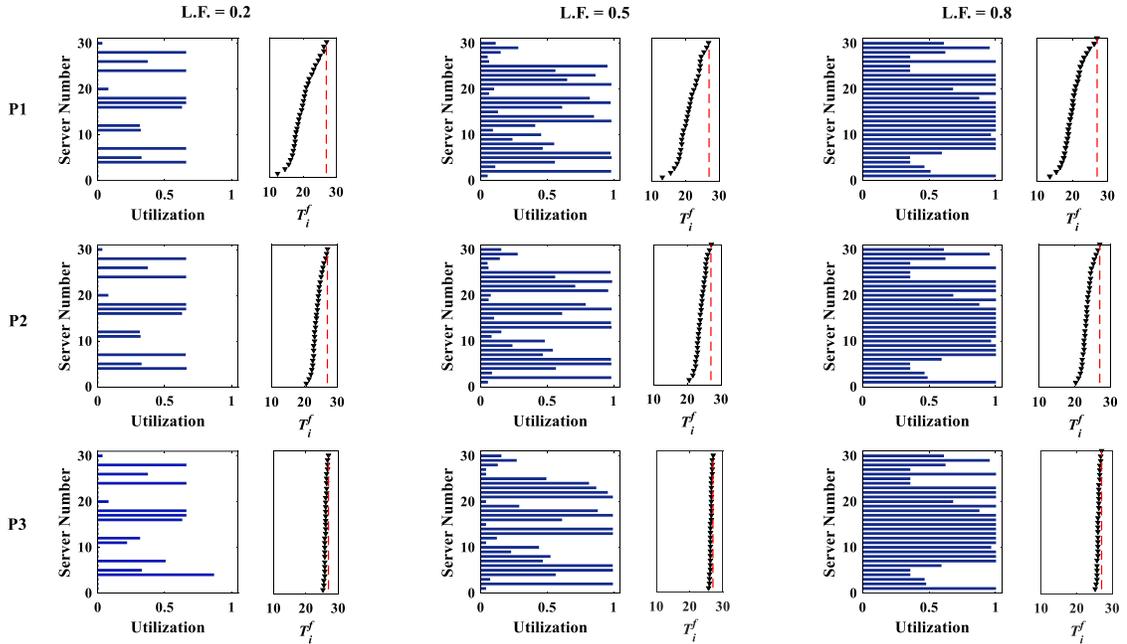


Figure 5.12: Utilization and temperature distributions for PUE and η_{2nd} aware multi-objective workload assignments. Nine salient points for the different LF marked in Figure 5.11 are considered for this diagram. The red dotted line corresponds to the maximum allowable server inlet temperature of 27 °C. The figure should be interpreted in accordance with Table 5.7.

By comparing point P2 marked on different LF curves in Figure 5.12, increasing LF does not significantly change the temperature distribution but significantly alters the workload assignment. Points P1 and P3 correspond to extrema where PUE is minimized and η_{2nd} is maximized, respectively. Optimizations 1 and 2 concentrate the IT load near

the RMCU for all values of LF . In contrast, the third optimization problem distributes the workload across the rack in a non-uniform manner so that a specific vertical temperature gradient is maintained while accounting for the tradeoff for two different objectives. Obtaining such a Pareto optimal front for a wide range of LF values provides an operational regime diagram (Figure 5.11) and helps realize potential savings across the objectives.

Table 5.7: Salient points for $LF = 0.2, 0.5$ and 0.8 obtained from Figure 5.11

LF	Point	\dot{Q}_a (CFM)	$T_{c,w}$ (°C)	η_{2nd}	PUE	Variance in cold aisle temperature distribution
0.2	1	515	10.1	12.0	1.58	12.8
	2	685	18.5	21.8	1.67	2.4
	3	839	21.5	28.5	1.78	0.1
0.5	1	531	10.6	12.9	1.51	10.3
	2	689	18.0	21.5	1.58	2.2
	3	850	22.0	28.3	1.68	0.1
0.8	1	611	10.8	16.9	1.49	8.9
	2	725	17.1	22.4	1.53	2.4
	3	850	22.0	27.5	1.60	0.1

5.4.6 Normalization of the Pareto front

To generalize the results for a single rack DC and demonstrate their applicability for multi-rack systems, the solutions obtained from multi-objective optimization are normalized. The baseline results obtained by solving single-objective problems at different LF are described in Figure 5.10. By fitting trend lines to each curve, the minima and maxima for PUE and η_{2nd} are obtained as functions of $0 \leq LF \leq 1$. The empirically obtained equations are as follows.

$$[PUE]_{min} = 1.64 - 0.39(LF) + 0.18(LF)^2, \quad (5.26)$$

$$[PUE]_{max} = 1.77 - 0.43(LF) + 0.15(LF)^2, \quad (5.27)$$

$$[\eta_{2nd}]_{min} = 5.27 + 53.41(LF) - 146.66(LF)^2 + 117.28(LF)^3, \text{ and} \quad (5.28)$$

$$[\eta_{2nd}]_{max} = 30.43 - 1.52(LF). \quad (5.29)$$

Therefore, the horizontal axis of the dimensionless Pareto front can be represented as,

$$PUE^* = \frac{[PUE]_i - [PUE]_{min}}{[PUE]_{max} - [PUE]_{min}}, \quad (5.30)$$

and the vertical axis is normalized as,

$$\eta_{2nd}^* = \frac{[\eta_{2nd}]_i - [\eta_{2nd}]_{min}}{[\eta_{2nd}]_{max} - [\eta_{2nd}]_{min}}, \quad (5.31)$$

where $[PUE]_i$ and $[\eta_{2nd}]_i$ represents the i^{th} point of the Pareto front.

Figure 5.13 represents the dimensionless variation of PUE vs. η_{2nd} where the effect of LF is normalized using Eqs. (5.26)-(5.31). There is a proportionality tradeoff across PUE^* and η_{2nd}^* where maximizing η_{2nd}^* leads to a detrimental effect on PUE^* . Thus, an intermediate value lying on the dimensionless Pareto front should be maintained during operation as per the choice of the DC administrator. The proposed normalization approach enhances the applicability of the results obtained for a single-rack system for a generalized multi-rack homogeneous DC providing better control over cooling-aware workload scheduling while ensuring the thermal reliability of the system.

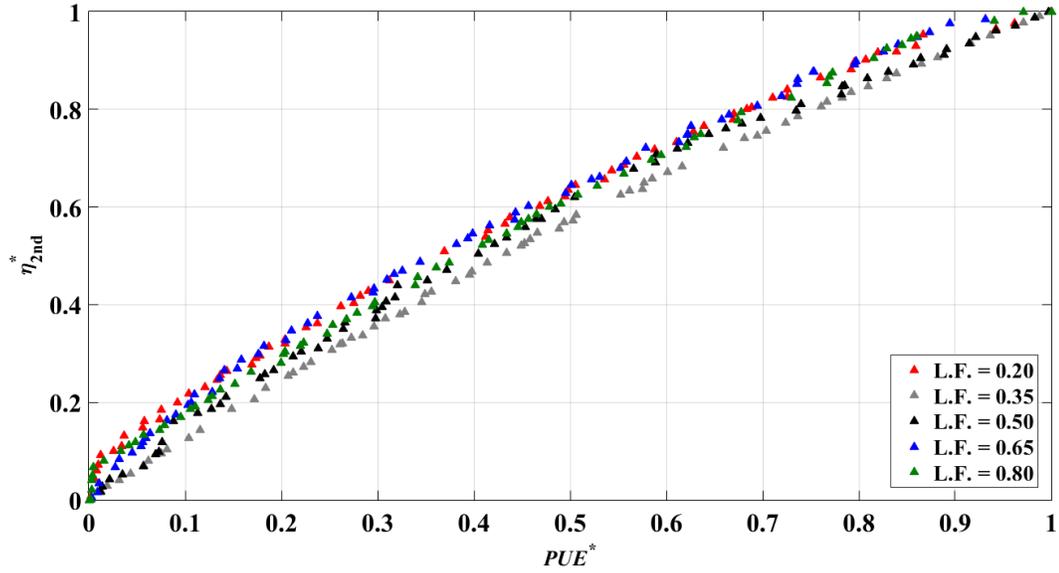


Figure 5.13: Dimensionless Pareto front of PUE and η_{2nd} obtained by adopting the proposed normalization procedure. The normalization procedure nullifies the effect of LF .

5.5 Conclusion

We have developed a novel approach to perform energy- and exergy-aware workload scheduling for air-cooled DCs with RMCUs. The methodology combines (1) a low complexity zonal temperature prediction model, (2) energy and exergy formulation, and (3) an optimization solver resulting in a facile tool for DC workload assignment. In contrast to existing thermal- and energy-aware IT load approaches, the irreversibility-based method opens up avenues for future research. The approach simultaneously regulates salient decision variables of the cooling system and the spatial workload distribution.

Important aspects of the investigation include elucidation of the impact on the algorithm of (1) altering cooling system variables and workload distribution on the energy and exergy efficiencies, (2) additional benefits obtained by considering the exergy efficiency, (3) tradeoffs across the three optimization problems, and (4) the influence of

the thermal reliability constraint, IT load factor and workload distribution. Salient findings are:

- Simultaneously increasing $T_{c,w}$ and \dot{Q}_a while maintaining $T_g \leq 27^\circ\text{C}$ increases the temperature uniformity and mean air temperature in the cold aisle.
- Populating the workload closer to the cooling unit lowers the value of \dot{Q}_a and reduces the power consumption of the fans.
- An increment in the dimensionless cooling parameter $\theta \times \beta$ improves the exergy efficiency at the cost of increasing the *PUE*. Therefore, the two desired values of the objectives cannot be achieved simultaneously.
- The *PUE* minimization problem does not necessarily minimize the overcooling of the servers in the vicinity of the RMCU. On the other hand, the exergy optimization problem lowers the extent of overcooling by increasing both $T_{c,w}$ and \dot{Q}_a .
- The η_{2nd} maximization problem results in an exergy efficiency of approximately 30% with a $\sim 7\%$ sacrifice on the optimal *PUE*.
- The multi-objective optimization problem distributes the workload across the IT rack, whereas the single-objective optimization problem assigns workload closest to the cooling unit.

The dimensionless Pareto front enhances the applicability of the results obtained for a single rack modular DC towards a multi-rack system.

5.6 References

- [1] W. Van Heddeghem, S. Lambert, B. Lannoo, D. Colle, M. Pickavet, and P. Demeester, "Trends in worldwide ICT electricity consumption from 2007 to 2012," *Computer Communications*, vol. 50, pp. 64-76, 2014.
- [2] S. V. Garimella, T. Persoons, J. Weibel, and L.-T. Yeh, "Technological drivers in data centers and telecom systems: Multiscale thermal, electrical, and energy management," *Applied energy*, vol. 107, pp. 66-80, 2013.
- [3] S. Herbert and D. Marculescu, "Analysis of dynamic voltage/frequency scaling in chip-multiprocessors," in *Proceedings of the 2007 international symposium on Low power electronics and design (ISLPED'07)*, 2007, pp. 38-43.
- [4] J. C. Smolens, R. T. Golla, and M. B. Smittle, "Dynamic mitigation of thread hogs on a threaded processor," ed: Google Patents, 2013.
- [5] J.-G. Park, C.-Y. Hsieh, N. Dutt, and S.-S. Lim, "Co-Cap: energy-efficient cooperative CPU-GPU frequency capping for mobile games," in *Proceedings of the 31st Annual ACM Symposium on Applied Computing*, 2016, pp. 1717-1723.
- [6] S.-T. Kung, C.-C. Cheng, C.-c. Liu, and Y.-c. Chen, "Dynamic power saving by monitoring CPU utilization," ed: Google Patents, 2003.
- [7] P. Padala, X. Zhu, Z. Wang, S. Singhal, and K. G. Shin, "Performance evaluation of virtualization technologies for server consolidation," *HP Labs Tec. Report*, vol. 137, 2007.

- [8] H. Moazamigoodarzi, P. J. Tsai, S. Pal, S. Ghosh, and I. K. Puri, "Influence of cooling architecture on data center power consumption," *Energy*, vol. 183, pp. 525-535, 2019.
- [9] J. D. Moore, J. S. Chase, P. Ranganathan, and R. K. Sharma, "Making Scheduling Cool": Temperature-Aware Workload Placement in Data Centers," in *USENIX annual technical conference, General Track*, 2005, pp. 61-75.
- [10] S. Zimmermann, I. Meijer, M. K. Tiwari, S. Paredes, B. Michel, and D. Poulikakos, "Aquasar: A hot water cooled data center with direct energy reuse," *Energy*, vol. 43, pp. 237-245, 2012.
- [11] K. Ebrahimi, G. F. Jones, and 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*, vol. 31, pp. 622-638, 2014.
- [12] S. Li, T. Abdelzaher, and M. Yuan, "Tapa: Temperature aware power allocation in data center with map-reduce," in *2011 international green computing conference and workshops*, 2011, pp. 1-8.
- [13] Y. Bai and L. Gu, "Chip temperature-based workload allocation for holistic power minimization in air-cooled data center," *Energies*, vol. 10, p. 2123, 2017.
- [14] R. K. Sharma, C. E. Bash, C. D. Patel, R. J. Friedrich, and J. S. Chase, "Balance of power: Dynamic thermal management for internet data centers," *IEEE Internet Computing*, vol. 9, pp. 42-49, 2005.

- [15] M. T. Chaudhry, T. Ling, S. A. Hussain, and A. Manzoor, "Minimizing thermal stress for data center servers through thermal-aware relocation," *The Scientific World Journal*, vol. 2014, 2014.
- [16] Q. Tang, S. K. S. Gupta, and G. Varsamopoulos, "Energy-efficient thermal-aware task scheduling for homogeneous high-performance computing data centers: A cyber-physical approach," *IEEE Transactions on Parallel and Distributed Systems*, vol. 19, pp. 1458-1472, 2008.
- [17] S. M. M. Nejad, G. Badawy, and D. G. Down, "EAWA: Energy-aware workload assignment in data centers," in *2018 International Conference on High Performance Computing & Simulation (HPCS)*, 2018, pp. 260-267.
- [18] X. Li, X. Jiang, P. Garraghan, and Z. Wu, "Holistic energy and failure aware workload scheduling in Cloud datacenters," *Future Generation Computer Systems*, vol. 78, pp. 887-900, 2018.
- [19] A. A. Khan, M. Zakarya, and R. Khan, "Energy-aware dynamic resource management in elastic cloud datacenters," *Simulation modelling practice and theory*, vol. 92, pp. 82-99, 2019.
- [20] D. S. Chisca, I. Castineiras, D. Mehta, and B. OSullivan, "On energy-and cooling-aware data centre workload management," in *2015 15th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing*, 2015, pp. 1111-1114.
- [21] S. MirhoseiniNejad, H. Moazamigoodarzi, G. Badawy, and D. G. Down, "Joint data center cooling and workload management: A thermal-aware approach," *Future Generation Computer Systems*, vol. 104, pp. 174-186, 2020.

- [22] A. J. Shah, V. P. Carey, C. E. Bash, and C. D. Patel, "Exergy-based optimization strategies for multi-component data center thermal management: Part II— application and validation," in *International Electronic Packaging Technical Conference and Exhibition*, 2005, pp. 215-224.
- [23] A. J. Shah, V. P. Carey, C. E. Bash, and C. D. Patel, "Exergy-based optimization strategies for multi-component data center thermal management: Part I—analysis," in *International Electronic Packaging Technical Conference and Exhibition*, 2005, pp. 205-213.
- [24] A. J. Shah, V. P. Carey, C. E. Bash, and C. D. Patel, "Exergy analysis of data center thermal management systems," *Journal of Heat Transfer*, vol. 130, 2008.
- [25] A. J. Shah, V. P. Carey, C. E. Bash, C. D. Patel, and R. K. Sharma, "Exergy Analysis of Data Center Thermal Management Systems," in *Energy Efficient Thermal Management of Data Centers*, Y. Joshi and P. Kumar, Eds., ed Boston, MA: Springer US, 2012, pp. 383-446.
- [26] A. J. Díaz, R. Cáceres, J. M. Cardemil, and L. Silva-Llanca, "Energy and exergy assessment in a perimeter cooled data center: The value of second law efficiency," *Applied Thermal Engineering*, vol. 124, pp. 820-830, 2017.
- [27] L. Silva-Llanca, A. Ortega, K. Fouladi, M. del Valle, and V. Sundaralingam, "Determining wasted energy in the airside of a perimeter-cooled data center via direct computation of the Exergy Destruction," *Applied Energy*, vol. 213, pp. 235-246, 2018.

- [28] L. Silva-Llanca, M. Del Valle, A. Ortega, and A. J. Díaz, "Cooling effectiveness of a data center room under overhead airflow via entropy generation assessment in transient scenarios," *Entropy*, vol. 21, p. 98, 2019.
- [29] A. Bhalerao, K. Fouladi, L. Silva-Llanca, and A. P. Wemhoff, "Rapid prediction of exergy destruction in data centers due to airflow mixing," *Numerical Heat Transfer, Part A: Applications*, vol. 70, pp. 48-63, 2016.
- [30] K. Fouladi, A. P. Wemhoff, L. Silva-Llanca, K. Abbasi, and A. Ortega, "Optimization of data center cooling efficiency using reduced order flow modeling within a flow network modeling approach," *Applied Thermal Engineering*, vol. 124, pp. 929-939, 2017.
- [31] R. Gupta, S. Asgari, H. Moazamigoodarzi, S. Pal, and I. K. Puri, "Cooling architecture selection for air-cooled Data Centers by minimizing exergy destruction," *Energy*, vol. 201, p. 117625, 2020.
- [32] H. Moazamigoodarzi, S. Pal, S. Ghosh, and I. K. Puri, "Real-time temperature predictions in it server enclosures," *International Journal of Heat and Mass Transfer*, vol. 127, pp. 890-900, 2018.
- [33] S. Pal, H. Moazami-goodarzi, S. Ghosh, K. Thorn, and I. K. Puri, "Plate-fin heat exchanger suitable for rack-mountable cooling unit," ed: Google Patents, 2019.
- [34] H. Moazamigoodarzi, S. Pal, D. Down, M. Esmalifalak, and I. K. Puri, "Performance of a rack mountable cooling unit in an IT server enclosure," *Thermal Science and Engineering Progress*, vol. 17, p. 100395, 2020.

- [35] H. Moazamigoodarzi, R. Gupta, S. Pal, P. J. Tsai, S. Ghosh, and I. K. Puri, "Modeling temperature distribution and power consumption in IT server enclosures with row-based cooling architectures," *Applied Energy*, vol. 261, p. 114355, 2020.
- [36] M. Kheradmandi, D. G. Down, and H. Moazamigoodarzi, "Energy-Efficient Data-Based Zonal Control of Temperature for Data Centers," in *IGSC*, 2019, pp. 1-7.
- [37] J. W. VanGilder, Z. Pardey, C. Healey, and X. Zhang, "A compact server model for transient data center simulations," *ASHRAE Transactions*, vol. 119, pp. 358-370, 2013.
- [38] M. Dayarathna, Y. Wen, and R. Fan, "Data center energy consumption modeling: A survey," *IEEE Communications Surveys & Tutorials*, vol. 18, pp. 732-794, 2015.
- [39] C. Jin, X. Bai, C. Yang, W. Mao, and X. Xu, "A review of power consumption models of servers in data centers," *Applied Energy*, vol. 265, p. 114806, 2020.
- [40] F. Chen, J. Grundy, Y. Yang, J.-G. Schneider, and Q. He, "Experimental analysis of task-based energy consumption in cloud computing systems," in *Proceedings of the 4th acm/spec international conference on performance engineering*, 2013, pp. 295-306.
- [41] S.-W. Ham, M.-H. Kim, B.-N. Choi, and J.-W. Jeong, "Simplified server model to simulate data center cooling energy consumption," *Energy and Buildings*, vol. 86, pp. 328-339, 2015.
- [42] H. Packard, "Quickspecs: HP ProLiant DL360 Generation 5 (G5)," in *Product datasheet*, 27, ed, 2007, pp. 1-43.

- [43] J. M. Gordon, K. Ng, and H. Chua, "Optimizing chiller operation based on finite-time thermodynamics: universal modeling and experimental confirmation," *International Journal of Refrigeration*, vol. 20, pp. 191-200, 1997.
- [44] TRANE, "Air-cooled liquid chillers, 10 to 60 Tons," ed, 2004.
- [45] S. Denki, "San Ace 80," in *Product catalog*, ed, 2016.
- [46] D. L. Beaty, "Data center energy metric: Power Usage Effectiveness (PUE)," *Ashrae Journal*, vol. 55, pp. 61-63, 2013.
- [47] M. P. David, M. Iyengar, P. Parida, R. Simons, M. Schultz, M. Gaynes, *et al.*, "Experimental characterization of an energy efficient chiller-less data center test facility with warm water cooled servers," in *2012 28th Annual IEEE Semiconductor Thermal Measurement and Management Symposium (SEMI-THERM)*, 2012, pp. 232-237.
- [48] M. Ouchi, Y. Abe, M. Fukagaya, H. Ohta, Y. Shinmoto, M. Sato, *et al.*, "Thermal management systems for data centers with liquid cooling technique of CPU," in *13th InterSociety Conference on Thermal and Thermomechanical Phenomena in Electronic Systems*, 2012, pp. 790-798.
- [49] Y. Pan, R. Yin, and Z. Huang, "Energy modeling of two office buildings with data center for green building design," *Energy and Buildings*, vol. 40, pp. 1145-1152, 2008.
- [50] R. Steinbrecher and R. Schmidt, "Data center environments," *ASHRAE J*, vol. 42, p. e49, 2011.

- [51] M. Iyengar and R. Schmidt, "Analytical modeling for thermodynamic characterization of data center cooling systems," *Journal of Electronic Packaging*, vol. 131, 2009.
- [52] D. W. Demetriou, H. E. Khalifa, M. Iyengar, and R. R. Schmidt, "Development and experimental validation of a thermo-hydraulic model for data centers," *HVAC&R Research*, vol. 17, pp. 540-555, 2011.
- [53] D. W. Demetriou, "Thermally aware, energy-based techniques for improving data center energy efficiency," 2012.
- [54] K. Brown, W. Torell, and V. Avelar, "Choosing the optimal data center power density," *Schneider White, Boston, MA*, p. 2, 2014.

6 Energy, exergy, and computing efficiency based workload and cooling management for multi-rack data centers

This chapter is reproduced from *Energy, exergy, and computing efficiency based data center workload and cooling management*, *Applied Energy* (2021): 117050 (In press), **Rohit Gupta**, Sahar Asgari, Hosein Moazamigoodarzi, Douglas G. Down, and Ishwar K. Puri, <https://doi.org/10.1016/j.apenergy.2021.117050>, Copyright © 2021, with permission from Elsevier. The author of this thesis is the first author and the primary contributor to this manuscript.

6.1 Abstract

The rapidly rising computing workloads in data centers (DCs) have necessitated new approaches to ensure effective performance and resilience that minimize the associated cooling energy. The literature on thermally-aware workload management provides strategies to reduce this energy cost while typically ignoring the reduction in cooling capacity due to thermodynamic irreversibility and computing performance per unit energy consumption. Hence, we provide an approach that considers coefficient of performance COP_c , exergy efficiency η_{ex} , and a new metric, computing performance ratio CPR . In contrast to existing methods that consider one-dimensional workload distributions, the temperature predictions from a physics-based zonal model are used to optimize cooling for two-dimensional workload distributions in a multi-rack DC. The investigation reveals

physics associated with two-dimensional workload management for multi-rack DCs, provides a framework for trade-offs between COP_c , η_{ex} , and CPR , explains the influence of IT load factor LF on different objectives and describes how parameters obtained from single- and multi-objective problems can vary. Our findings show that COP_c , and η_{ex} can be improved by up to 20% and 8% by regulating the chilled water temperature and airflow setpoints while increasing the LF degrades the CPR by 7.5%. These results enable an extended approach for heterogeneous LF management in large-scale DCs.

6.2 Introduction

Data centers (DCs) and high-performance computing (HPC) clusters nowadays routinely process workloads associated with analytics and machine learning applications. For actionable use, high-performance computing infrastructure with low latencies is necessary. Since heat dissipation from DCs and HPC clusters has increased significantly with the use of multi-core processors and three-dimensional chip stacking, thermal and workload management requires adequate cooling at low cost. There are several thermal and workload management strategies to address operational expenditure [1, 2], chip thermal reliability [3], computing performance [4], latency [5], cooling architecture [6, 7], and waste heat harnessing [8-10].

Thermally-aware workload management and cooling is an integrated approach that considers interrelated issues, such as DC operational cost, chip thermal reliability, computing performance, and cooling architecture. The thermal reliability of computing chips is degraded because of (1) transient temperature fluctuations that diminish their lifetimes and (2) steady-state operating temperatures that rise above the critical allowable

chip temperature and induce thermal shutdown. These thermal issues can be addressed by utilizing chip temperature-aware and cold aisle temperature-aware workload management approaches [11]. Other workload scheduling approaches address thermal guidelines and operational expenditure by minimizing heat recirculation across the cold and hot air streams in a DC [12], diminishing hot spots among neighboring servers [13], and lowering temperature nonuniformities in server exhausts [14, 15] (see Table 6.1).

Although it is possible to ensure thermal reliability while minimizing DC operational expenditure, existing temperature-based methodologies do not address the reduction in cooling capacity due to thermodynamic irreversibilities that result from uneven workload distributions and cooling parameters. This deficiency can be addressed by approaches for exergy-based cooling management that minimize heat recirculation and bypass [16, 17], identify component level inefficiencies [6, 18], and optimize cooling parameters [19, 20]. However, there are few similar applications for workload management.

Recent investigations demonstrate the benefits of an exergy-based approach to select cooling architecture [6] and manage workload in air-cooled DCs [21]. The former advocates for the benefits of a modular DC cooling architecture over legacy raised-floor DCs, showing a 23% increase in the effective cooling capacity. The latter optimizes energy and exergy efficiency jointly for a single-rack modular air-cooled DC equipped with a rack-mountable cooling unit while simultaneously performing one-dimensional workload management and regulating cooling parameter setpoints. A comprehensive review of state of the art in thermal-aware workload management literature is presented in Table 6.1. Some

important aspects highlighted in this table are the number and type of objectives considered in the workload management framework, choice of the thermal model, type of the DC cooling architecture investigated, and whether the approach is multi-objective in nature to examine trade-offs between metrics.

Scaling up from a single-rack DC to a multi-rack system dramatically alters the airflow and cooling characteristics. Because there are shared hot and cold airspaces in a multi-rack infrastructure, one-dimensional workload distributions obtained for a single-rack system cannot be scaled up for multi-rack DCs [21, 22]. Thus, workload distributions that vary in two dimensions must be obtained. Previous investigations have also not considered computing performance, i.e., Floating Point Operations Per Second (FLOPs) per Watt of energy consumption [23, 24], simultaneously with energy and exergy efficiency. If an air-cooled DC is managed solely based on energy and exergy efficiencies, computing performance can be drastically degraded as electronic chips become overheated. The actual FLOPs per Watt $\eta_{c,a}$ is determined by performing LINPACK benchmark tests on a processor of interest [23]. Further, this quantity can be normalized using the maximum achievable FLOPs per Watt $\eta_{c,i}$, which helps to define the dimensionless metric computing performance ratio CPR . The value of $\eta_{c,a}$ is regulated by changing the (1) CPU frequency through over- or under- clocking, (2) voltage, and (3) temperature [25, 26]. Although simultaneous optimization of these parameters results in a global optimal value of CPR , we only focus on thermal performance as per our objective.

We address the shortcomings of prior workload management approaches by proposing a novel methodology to maximize (1) energy, (2) exergy, and (3) computing

efficiencies for a five-rack air-cooled DC equipped with two in-row cooling (IRC) units [22]. Energy and flow interactions are considered across DC racks to obtain spatially two-dimensional temperature distributions, which are used to solve a multi-objective genetic algorithm (MOGA)-based optimization problem to optimize two-dimensional workload distributions and cooling parameters, such as chilled water temperature and airflows of the IRC units.

Table 6.1: Relevant state-of-the-art works in the area of thermal-aware workload and cooling management in DCs.

Ref	Actual DC	Objectives or goals	Type of thermal model	Auxiliary sub-models	Optimization
[27]	No	<ul style="list-style-type: none"> Maximizing utilization of solar energy 	Heat recirculation matrix (HRM)	<ul style="list-style-type: none"> Regression-based non-linear cooling energy model Linear server energy model Neural network-based solar energy model 	Single-objective optimization with linear constraints
[11]	Raised-floor DC	<ul style="list-style-type: none"> Minimizing total energy consumption (including servers) 	HRM	<ul style="list-style-type: none"> Regression-based non-linear cooling energy model Non-linear server energy model Non-linear CPU temperature model Proper orthogonal decomposition (POD) based 	Single-objective optimization with non-linear constraints

				cold chamber temperature model	
[13]	Raised-floor DC	<ul style="list-style-type: none"> Minimizing heat recirculation by assigning workloads according to cold chamber temperature and load of neighboring servers 	Computational fluid dynamics (CFD) simulation	<ul style="list-style-type: none"> Linear server energy model Regression-based cold chamber temperature model 	No
[28]	No	<ul style="list-style-type: none"> Minimize hotspot temperature 	HRM	<ul style="list-style-type: none"> Linear server energy model Regression-based cooling energy model Job processing time models 	Single-objective optimization with non-linear constraints
[29]	Raised-floor DC	<ul style="list-style-type: none"> Minimizing peak temperature of servers Minimizing job processing time of servers Minimizing CO₂ emission 	CFD simulation	<ul style="list-style-type: none"> Thermal resistance capacitance-based CPU temperature model CFD data-driven cold aisle temperature predictor Job processing time models 	Single-objective optimization with non-linear constraints
[12]	Raised-floor DC	<ul style="list-style-type: none"> Minimizing hot spot air temperature in cold chamber 	HRM	<ul style="list-style-type: none"> Linear server energy model Non-linear regression-based cooling energy model 	Single-objective optimization with non-linear constraints

[29]	Raised-floor DC	<ul style="list-style-type: none"> Minimizing total cooling energy Minimize CPU temperature 	CFD simulation	<ul style="list-style-type: none"> Regression-based CRAC energy model Thermal resistance capacitance-based CPU temperature model CFD data-driven cold aisle temperature predictor 	Single-objective optimization with non-linear constraints
[15]	No	<ul style="list-style-type: none"> Minimizing hotspot temperature in cold chamber 	Regression-based	<ul style="list-style-type: none"> Server thermal stress model 	No
[30]	Raised-floor DC	<ul style="list-style-type: none"> Minimizing total energy consumption (including servers) 	Regression-based	<ul style="list-style-type: none"> Non-linear regression-based server energy model Non-linear regression-based cooling energy model 	Single-objective optimization-based model predictive control
[31]	Raised-floor DC	<ul style="list-style-type: none"> Maximizing utilization of solar photovoltaic energy Maximizing utilization of free cooling 	Regression-based	<ul style="list-style-type: none"> Regression-based cooling energy model Solar photovoltaic power forecasting model Free cooling system model Operational cost model 	Single-objective optimization with linear and non-linear constraints
[14]	Raised-floor DC	<ul style="list-style-type: none"> Maintain uniformity in the server 	No	No	No

		<p>exhaust temperatures</p> <ul style="list-style-type: none"> • Reduce overcooling of the cold aisle 			
[20]	Raised-floor DC	<ul style="list-style-type: none"> • Minimizing energy consumption of cooling cycle • Minimizing exergy destruction in DC (including servers) 	CFD simulation data-driven POD model	<ul style="list-style-type: none"> • Energy consumption and exergy destruction models of cooling cycle components (server, heat exchanger, fan, pump, chiller, and cooling tower) 	Unconstrained single-objective optimization
[1]	single-rack DC with RMCU	<ul style="list-style-type: none"> • Minimizing cooling energy consumption 	FNR-based zonal model	<ul style="list-style-type: none"> • Linear server energy model • Non-linear chiller power model • Non-linear fan power model 	Single-objective optimization problem with both linear and non-linear constraints
[21]	single-rack DC with RMCU	<ul style="list-style-type: none"> • Minimizing cooling energy consumption • Maximizing exergy efficiency of cooling cycle 	FNR-based zonal model	<ul style="list-style-type: none"> • Linear server energy model • Non-linear chiller power model • Linear fan model • Exergy destruction model of each component 	Dual-objective optimization problem with both linear and non-linear constraints
Present work	Multi-rack DC with IRC units	<ul style="list-style-type: none"> • Minimize cooling energy consumption • Maximize exergy efficiency of 	FNR-based zonal model	<ul style="list-style-type: none"> • Linear server energy model • Non-linear CPU temperature model 	Tri-objective optimization problem with linear and non-linear constraints

		cooling cycle <ul style="list-style-type: none"> Maximize computing performance per unit energy consumption 		<ul style="list-style-type: none"> Non-linear chiller power model Non-linear fan power model Exergy destruction model of each DC component Non-linear computing performance per watt model 	
--	--	---	--	--	--

As observed from Table 6.1, thermally-aware workload distribution algorithms require an accurate dynamic thermal model for the DC. Several options are available, such as (a) data-driven black-box models [32], (b) data-driven hybrid gray-box models [33, 34], (c) computational fluid dynamics (CFD) simulation-based heat recirculation matrix (HRM) approaches [12], and (d) physics-based zonal models [22, 35]. The first three require significant training data and computational resources and are the basis of prior workload scheduling algorithms. To circumvent this problem, our framework utilizes an experimentally validated, training-free, physics-based, spatiotemporal zonal model to predict temperatures in a five-rack DC [22]. This physics-based approach eliminates extrapolative temperature predictions, reduces computational time, and improves the accuracy of temperature predictions for new scenarios.

To the best of our knowledge, this is the first work that develops a comprehensive workload and cooling management framework for multi-rack DCs cooled by IRC units. Considering the energy, exergy, and computing efficiency along with thermal reliability

constraints allows for the exploration of insightful operational trade-offs. The proposed multi-objective optimization problem not only reveals the cooling and spatial workload dynamics but also provides multi-objective trade-off diagrams applicable for real-world DC applications.

With this background information, our sixfold contributions are:

- Proposing a new metric to describe computing performance degradation with rising CPU temperatures or thermal throttling. It has been observed in this work that with an elevation of CPU temperatures, the computing performance ratio can degrade by up to 7.5%.
- Explaining the trade-offs between energy and exergy-based cooling metrics and the computing performance metric. We solve three single- and one multi-objective optimization problems to better understand the nature of the objective function in decision making.
- Deriving a comprehensive formulation that couples our previous flow network-based temperature prediction model [22] for a five-rack air-cooled DC equipped with IRC units with energy, exergy, and computing efficiency models.
- Understanding of the relationship between workload distribution, cooling cycle operation, and thermal characteristics of a DC with shared hot and cold chambers.
- Explaining the influence of the IT load factor for different optimization scenarios. We find that the load factor significantly alters the energy and computing efficiencies while the optimal value of exergy efficiency remains nearly unaltered by changing load factor.

- Utilizing the results for the five-rack case study to develop an extended scalable workload management approach for larger DCs.

The remaining sections are organized as follows. Section 2 develops a series of optimization problems with different objectives informed by the flow network-based temperature prediction model. Section 3 includes the results and physical insights obtained by solving the optimization problems. Finally, Section 4 provides concluding remarks and highlights salient findings.

6.3 Methodology

6.3.1 Physical layout and thermal model of the row-based cooling infrastructure

The multi-rack DC architectural layout with IRC units is described in Figure 6.1 [22]. The DC cooling infrastructure consists of five IT racks, each of them containing 15 2U servers (where 1U = 4.4 cm). The racks are stacked within an enclosure and share a common cold and hot chamber at the front and back of the IT racks, respectively. The cold chamber and hot chamber are separated by high-density air blocking brushes to prevent air from mixing. Two IRC units manufactured by RITTAL [36] supply the required cold air from two sides of the IT enclosure. Each IRC unit consists of an air-water fin-tube heat exchanger, three centrifugal fans, and valve-based water flow control mechanism. A vapor compression refrigeration (VCR) chiller, equipped with an ambient air-cooled condenser [37], supplies the required chilled water to the IRC heat exchangers.

The thermal model adopts a zonal approach, which assumes that physical quantities inside each zone are spatially uniform [22]. This is a faster and reasonably accurate alternative to a full-field CFD simulation. A zone containing servers is modelled as a heat

source whose magnitude depends on the IT stress level or utilization. Cooling unit fan zones supply predefined airflows at a specified temperature.

There are several types of zones inside the DC enclosure: (1) zones in front of the IRC units, i.e., the cold air supply zone, (2) zones at the back of the IRC units, i.e., warm air return zone, (3) cold air intake zones in front of the servers, (4) hot air exhaust zones at the back of servers, (5) each server itself, and (6) the IRC units. The cold chamber and hot chamber are partitioned into 25 zones of equal volume so that each consists of three 2U servers. The rationale for filling the IT racks in this manner is provided by the temperature prediction model, where each rack of 30U height must be filled (30U per rack = 5 zones per rack \times 3 servers per zone \times 2U server) [22, 35]. Additionally, there are five cold air supply and return zones on each side of the enclosure, as shown in Figure 6.1b. Temperature prediction using the zonal framework is a two-step process, where a flow network representation (FNR) is used to calculate the pressures and airflows of all the zones. Thereafter, the pressure-flow data are used in energy balance equations to determine the zonal temperatures.

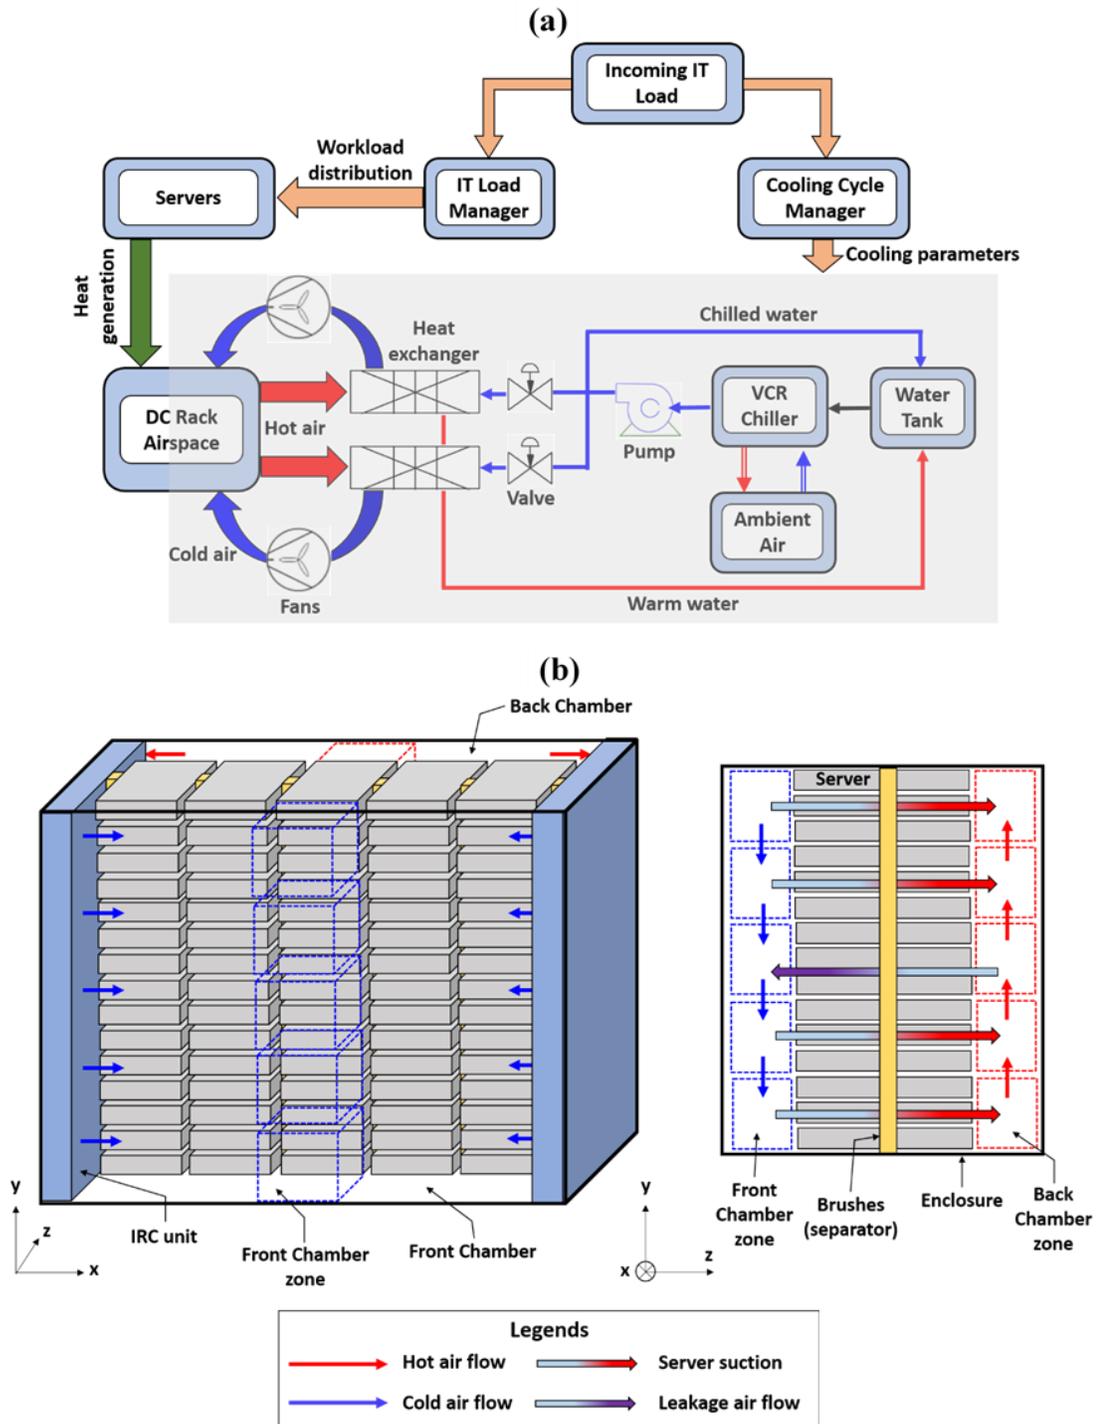


Figure 6.1: (a) Process flow layout across different components for the DC case study, and (b) representative schematic of the DC infrastructure equipped with five server racks and

two IRC units. The diagram on the left shows a three-dimensional representation of the DC, whereas the one on the right shows a sectional view representing salient airflows. A server rack is divided into five zones, where each zone consists of three 2U servers (1U = 4.4 cm). There are 35 zones (25 zones in front of servers and 5 zones in front of each IRC unit) in each of the front and back chambers.

6.3.1.1 Airflow estimation of front (cold) and back (hot) chambers

The flow-field inside an enclosed DC is pressure-driven [22, 38], enabling an FNR to calculate the pressure-flow characteristics. The FNR is based on our previous work for a five-rack DC with two IRC units, as shown in Figure 6.2 [22]. IRC zones are represented as sources of airflow \dot{Q}_a , whereas the server airflows are denoted by $\dot{Q}_{i,j}^z$. The horizontal flow transport resistances in the front and back chambers are R_H^f and R_H^b , respectively, whereas the vertical flow transport resistances for the front and back chambers are R_V^f and R_V^b . The high-density air blocking brushes are porous separation media enabling momentum and energy flux transport across the two chambers, which can be represented by a flow resistance R_{br} , values for which have been reported [22]. Since the enclosure is air-tight, there are no flow interactions across the rack and the room (ambient).

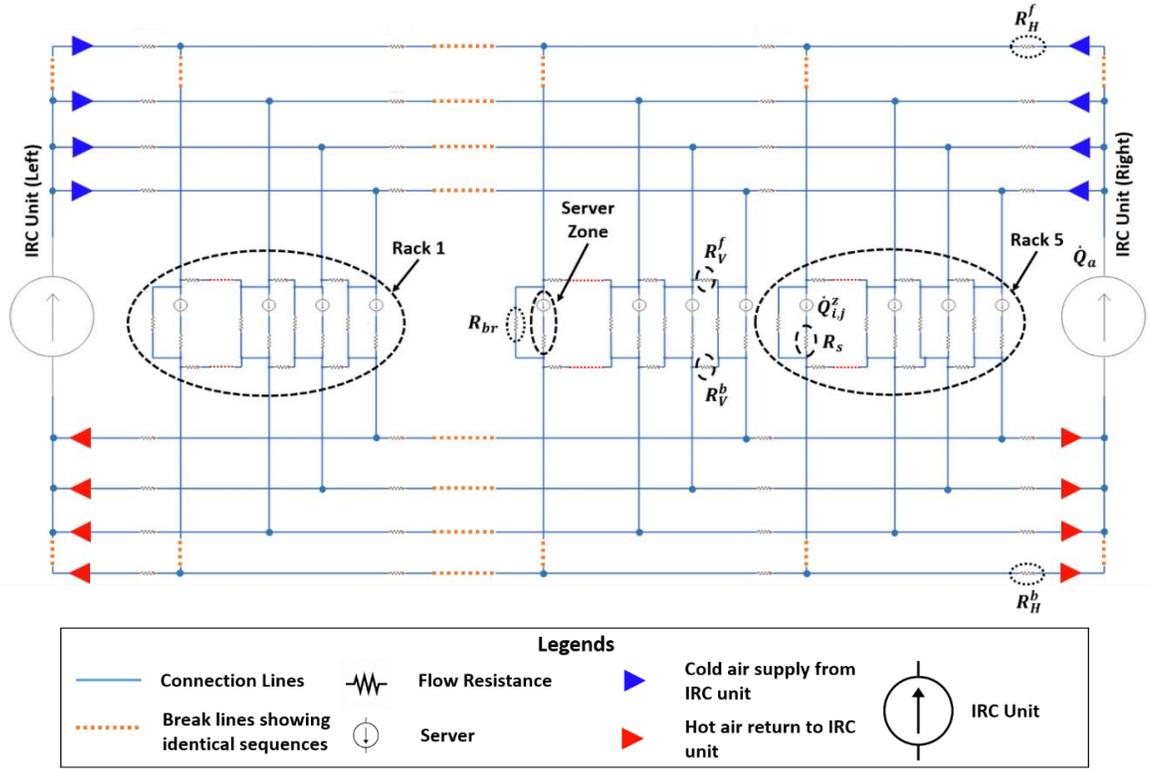


Figure 6.2: Generalized flow network representation (FNR) of the DC infrastructure for the case study with a total of 75 2U servers and 2 IRC units. The orange dotted lines represent identical sequences of flow resistances and servers in a simplified manner. Reproduced with permission from Ref. [22].

As a function of inlet air temperature, $T_{i,j}^f$, the airflow through a zone with N_s 2U servers are [35],

$$\dot{Q}_{i,j}^z = \begin{cases} N_s \times 0.01415 & \text{if } T_{i,j}^f \leq 25 \text{ }^\circ\text{C} \\ N_s \times [0.01415 + (T_{i,j}^f - 25) \times 0.00142] & \text{if } 25 < T_{i,j}^f < 35 \text{ }^\circ\text{C} \end{cases} \quad (6.1)$$

For our system $N_s = 3$ as there are three 2U servers in each zone of an IT rack.

With flow resistances known, mass balance equations for the front chamber, back chamber, and heat-generating server zones can be formulated. The mass balances for zones in contact with active servers in the front and back chambers are expressed by Eqs. (6.2)

and (6.3). The mass balances for the zones in an IRC cooling unit supply and return are calculated using Eqs. (6.4) and (6.5).

$$\alpha^f [P_{i+1,j}^f + P_{i-1,j}^f] + \beta^f [P_{i,j+1}^f + P_{i,j-1}^f] + \gamma P_{i,j}^b - (2\alpha^f + 2\beta^f + \gamma)P_{i,j}^f - \dot{Q}_{i,j}^z = 0, \quad (6.2)$$

$$\alpha^b [P_{i+1,j}^b + P_{i-1,j}^b] + \beta^b [P_{i,j+1}^b + P_{i,j-1}^b] + \gamma P_{i,j}^f - (2\alpha^b + 2\beta^b + \gamma)P_{i,j}^b + \dot{Q}_{i,j}^z = 0, \quad (6.3)$$

$$\alpha^f [P_{i+1,j}^f - P_{i,j}^f] - \frac{\dot{Q}_a}{N_z N_{cu}} = 0, \quad (6.4)$$

$$\alpha^b [P_{i+1,j}^b - P_{i,j}^b] + \frac{\dot{Q}_a}{N_z N_{cu}} = 0, \quad (6.5)$$

where $\alpha^f = 1/R_H^f$, $\alpha^b = 1/R_H^b$, $\beta^f = 1/R_V^f$, $\beta^b = 1/R_V^b$, $\gamma = 1/R_{br}$, \dot{Q}_a the total airflow setpoint prescribed by the DC control system (split equally across two IRC units), $N_{cu} = 2$ is the number of IRC units, $N_z = 5$ the number of zones inside a rack, and P the pressure corresponding to the zone of interest. Equations (6.2)-(6.5) constitute a system of linear equations whose solution yields the spatial pressure distribution across different zones in the front and back chambers (see Appendix of Chapter 6 for a detailed description).

6.3.1.2 Temperature estimation of front (cold) and back (hot) chambers

With the pressure distribution determined using FNR, the temperatures in the front and back chambers can be calculated. The transient energy balance for a server zone based on the forward time-marching method is,

$$\frac{X_s}{2} \left(\frac{T_{i,j}^e|_t - T_{i,j}^e|_{t-\Delta t}}{\Delta t} + \frac{T_{i,j}^f|_t - T_{i,j}^f|_{t-\Delta t}}{\Delta t} \right) = \rho_a C_{p,a} \dot{Q}_{i,j}^z (T_{i,j}^f|_{t-\Delta t} - T_{i,j}^e|_{t-\Delta t}) + \dot{P}_{i,j}^z, \quad (6.6)$$

where X_s is the thermal mass of a 2U server (available from the literature) [39], $T_{i,j}^e$ and $T_{i,j}^f$ are the server exhaust and inlet air temperatures, ρ_a the density of air, $C_{p,a}$ the specific heat capacity of air, Δt the time step, and $\dot{P}_{i,j}^z$ the power consumption of the server zone with horizontal rack index i , and vertical zone index j . We note that a server's power consumption depends on its utilization, which is one of our decision variables for which the appropriate function is provided in Eq. (6.11).

The energy balance equations for the waterside and airside of the heat exchanger inside each IRC unit are expressed by Eqs. (6.7) and (8), respectively.

$$\frac{X_w}{2} \left(\frac{T_{h,w}^t - T_{h,w}^{t-\Delta t}}{\Delta t} \right) = \frac{UA}{2} (T_{h,a}^{t-\Delta t} + T_{c,a}^{t-\Delta t} - T_{c,w} - T_{h,w}^{t-\Delta t}) + \rho_w C_{p,w} \frac{\dot{Q}_w}{N_{cu}} (T_{c,w} - T_{h,w}^{t-\Delta t}), \quad (6.7)$$

$$\begin{aligned} \frac{X_a}{2} \left(\frac{T_{h,a}^t - T_{h,a}^{t-\Delta t}}{\Delta t} + \frac{T_{c,a}^t - T_{c,a}^{t-\Delta t}}{\Delta t} \right) \\ = -\frac{UA}{2} (T_{h,a}^{t-\Delta t} + T_{c,a}^{t-\Delta t} - T_{c,w} - T_{h,w}^{t-\Delta t}) + \rho_a C_{p,a} \frac{\dot{Q}_a}{N_{cu}} (T_{h,a}^{t-\Delta t} - T_{c,a}^{t-\Delta t}). \end{aligned} \quad (6.8)$$

Here $X_w = \rho_w C_{p,w} V_w$ and $X_a = \rho_a C_{p,a} V_a$ are the thermal masses of water and air inside the IRC unit, \dot{Q}_w the total water flow prescribed by the DC control system (distributed equally across two IRC units), $T_{h,w}$ and $T_{c,w}$ the hot and chilled water temperatures, $T_{h,a}$ and $T_{c,a}$ the hot air return and cold air supply temperatures, UA the product of the universal heat transfer coefficient and the contact area between the two interacting fluid media, i.e., air and water. The value of UA as a function of \dot{Q}_a and \dot{Q}_w is obtained from our previous work [22].

Table 6.2: Expressions for the zonal flow interaction terms $\Omega_{1 \rightarrow 6}^f$ and $\Omega_{1 \rightarrow 6}^b$ used in Eqs. (6.9) and (6.10), respectively.

$\Omega_{1 \rightarrow 6}^f$		$\Omega_{1 \rightarrow 6}^b$	
Ω_1^f (Horizontal flow transport in front chamber)		Ω_1^b (Horizontal flow transport in back chamber)	
$[P_{i+1,j}^f - P_{i,j}^f] \geq 0$	$\rho_a C_{p,a} \alpha^f (P_{i+1,j}^f - P_{i,j}^f) T_{i+1,j}^f$	$[P_{i+1,j}^b - P_{i,j}^b] \geq 0$	$\rho_a C_{p,a} \alpha^b (P_{i+1,j}^b - P_{i,j}^b) T_{i+1,j}^b$
$[P_{i+1,j}^f - P_{i,j}^f] < 0$	$\rho_a C_{p,a} \alpha^f (P_{i+1,j}^f - P_{i,j}^f) T_{i,j}^f$	$[P_{i+1,j}^b - P_{i,j}^b] < 0$	$\rho_a C_{p,a} \alpha^b (P_{i+1,j}^b - P_{i,j}^b) T_{i,j}^b$
Ω_2^f (Horizontal flow transport in front chamber)		Ω_2^b (Horizontal flow transport in back chamber)	
$[P_{i-1,j}^f - P_{i,j}^f] \geq 0$	$\rho_a C_{p,a} \alpha^f (P_{i-1,j}^f - P_{i,j}^f) T_{i-1,j}^f$	$[P_{i-1,j}^b - P_{i,j}^b] \geq 0$	$\rho_a C_{p,a} \alpha^b (P_{i-1,j}^b - P_{i,j}^b) T_{i-1,j}^b$
$[P_{i-1,j}^f - P_{i,j}^f] < 0$	$\rho_a C_{p,a} \alpha^f (P_{i-1,j}^f - P_{i,j}^f) T_{i,j}^f$	$[P_{i-1,j}^b - P_{i,j}^b] < 0$	$\rho_a C_{p,a} \alpha^b (P_{i-1,j}^b - P_{i,j}^b) T_{i,j}^b$
Ω_3^f (Vertical flow transport in front chamber)		Ω_3^b (Vertical flow transport in back chamber)	
$[P_{i,j+1}^f - P_{i,j}^f] \geq 0$	$\rho_a C_{p,a} \beta^f (P_{i,j+1}^f - P_{i,j}^f) T_{i,j+1}^f$	$[P_{i,j+1}^b - P_{i,j}^b] \geq 0$	$\rho_a C_{p,a} \beta^b (P_{i,j+1}^b - P_{i,j}^b) T_{i,j+1}^b$
$[P_{i,j+1}^f - P_{i,j}^f] < 0$	$\rho_a C_{p,a} \beta^f (P_{i,j+1}^f - P_{i,j}^f) T_{i,j}^f$	$[P_{i,j+1}^b - P_{i,j}^b] < 0$	$\rho_a C_{p,a} \beta^b (P_{i,j+1}^b - P_{i,j}^b) T_{i,j}^b$
Ω_4^f (Vertical flow transport in front chamber)		Ω_4^b (Vertical flow transport in back chamber)	
$[P_{i,j-1}^f - P_{i,j}^f] \geq 0$	$\rho_a C_{p,a} \beta^f (P_{i,j-1}^f - P_{i,j}^f) T_{i,j-1}^f$	$[P_{i,j-1}^b - P_{i,j}^b] \geq 0$	$\rho_a C_{p,a} \beta^b (P_{i,j-1}^b - P_{i,j}^b) T_{i,j-1}^b$
$[P_{i,j-1}^f - P_{i,j}^f] < 0$	$\rho_a C_{p,a} \beta^f (P_{i,j-1}^f - P_{i,j}^f) T_{i,j}^f$	$[P_{i,j-1}^b - P_{i,j}^b] < 0$	$\rho_a C_{p,a} \beta^b (P_{i,j-1}^b - P_{i,j}^b) T_{i,j}^b$
Ω_5^f (Flow transport through brushes in front chamber)		Ω_5^b (Flow transport through brushes in back chamber)	
$[P_{i,j}^f - P_{i,j}^f] \geq 0$	$\rho_a C_{p,a} \gamma (P_{i,j}^f - P_{i,j}^f) T_{i,j}^f$	$[P_{i,j}^b - P_{i,j}^b] \geq 0$	$\rho_a C_{p,a} \gamma (P_{i,j}^b - P_{i,j}^b) T_{i,j}^b$
$[P_{i,j}^f - P_{i,j}^f] < 0$	$\rho_a C_{p,a} \gamma (P_{i,j}^f - P_{i,j}^f) T_{i,j}^f$	$[P_{i,j}^b - P_{i,j}^b] < 0$	$\rho_a C_{p,a} \gamma (P_{i,j}^b - P_{i,j}^b) T_{i,j}^b$
Ω_6^f (Flow through active server)		Ω_6^b (Flow through active server)	
$-\rho_a C_{pa} \dot{Q}_{i,j}^s T_{i,j}^f$		$\rho_a C_{pa} \dot{Q}_{i,j}^s T_{i,j}^e$	

For the zones considered within the front and back chambers, the energy balance equations are,

$$X_z \left(\frac{T_{i,j}^f|_t - T_{i,j}^f|_{t-\Delta t}}{\Delta t} \right) = \Omega_1^f + \Omega_2^f + \Omega_3^f + \Omega_4^f + \Omega_5^f + \Omega_6^f, \quad (6.9)$$

$$X_z \left(\frac{T_{i,j}^b|_t - T_{i,j}^b|_{t-\Delta t}}{\Delta t} \right) = \Omega_1^b + \Omega_2^b + \Omega_3^b + \Omega_4^b + \Omega_5^b + \Omega_6^b, \quad (6.10)$$

where $X_z = \rho_a C_{p,a} V_z$ denotes the thermal mass of zones in the front and back chambers and Ω_1 through Ω_6 are zonal flow interaction terms calculated using the pressure distribution. The expressions of these terms for different types of zones are provided in Table 6.2. Equations (6.6)-(6.10) are solved by forward time marching to determine the temperature distributions for the zones in front and back chambers [22].

6.3.2 Objective 1: Energy consumption

The energy consumption in a DC consists of two components: (1) the servers and (2) the cooling cycle. Intuitively, the power consumption of a server is a composite function of its utilization and the inlet air temperature [40]. However, since changing the inlet air temperature by 40% alters the server energy consumption by only 1% [41], we assume that the energy consumption attributed to a server is solely a function of its utilization. The total IT energy consumption inside the DC can therefore be expressed as,

$$\dot{P}_{IT} = \sum_{j=1}^{N_z} \sum_{i=1}^{N_r} \dot{P}_{i,j}^z = \sum_{j=1}^{N_z} \sum_{i=1}^{N_r} N_s [A_1 + A_2 u_{i,j}^z], \quad (6.11)$$

where the number of racks $N_r = 5$, number of servers per zone $N_s = 3$, and the number of zones in a rack $N_z = 5$. The server power model constants $A_1 = 223.4$ and $A_2 = 154.5$ are obtained from the datasheet of an HP ProLiant DL360 G5 server with two Intel Xeon E5-2697 v3 processors [42]. We assume this particular server to be employed solely throughout our DC.

The energy consumption of the cooling cycle has three components: (1) fans, (2) VCR chiller, and (3) pumps. The DC architecture is equipped with a fixed speed pump with valve-based flowrate control. Since the third component is invariant across different

operational conditions and cannot be optimized, we only optimize the energy consumption of the fans inside the IRC units and the VCR chiller. The total energy consumption of the cooling cycle that can be optimized is,

$$\dot{P}_{cool} = \dot{P}_f + \dot{P}_{ch}. \quad (6.12)$$

The fan energy consumption for the RITTAL IRC units can be expressed as [22],

$$\dot{P}_f = N_{cu} \left[B_1 + B_2 \left(\frac{\dot{Q}_a}{N_{cu}} \right) + B_3 \left(\frac{\dot{Q}_a}{N_{cu}} \right)^2 \right], \quad (6.13)$$

where $B_1 = 480$, $B_2 = -3073$, and $B_3 = 6031$ are fan model constants [22]. We note that fan energy consumption models can be parabolic, cubic, or linear for different manufacturers and classes of fans.

While there are several options for calculating the energy consumption of the VCR chiller, we utilize the well-known Gordon-Ng universal chiller model [43] for its ease of use and low computational requirement, and high accuracy across a wide range of heat loads. The model has the form,

$$\dot{P}_{ch} = \frac{\dot{Q}_{ch} + C_1 T_{c,w} T_{amb} + C_2 (T_{amb} - T_{c,w})}{T_{c,w} - C_3 \dot{Q}_{ch}} \dot{Q}_{ch}, \quad (6.14)$$

where \dot{Q}_{ch} is the heat load on the chiller evaporator, i.e., the server heat load, $T_{c,w}$ the chilled water setpoint prescribed by the DC control system, T_{amb} the ambient air temperature entering the chiller condenser, and \dot{P}_{ch} the power consumption of the chiller. All temperatures in Eq. (6.14) are expressed in units of K, whereas \dot{P}_{ch} and \dot{Q}_{ch} are in kW. To evaluate the model constants C_1 to C_3 , performance data of the commercially available TRANE CGA120 10 Tons chiller is used [37]. The constants evaluated using multivariate

linear regression are $C_1 = 0.0092$, $C_2 = 8.953$, and $C_3 = 0.649$. The value of T_{amb} is taken as 297.04 K from the chiller performance datasheet.

Finally, the performance of the DC cooling system from the 1st law of thermodynamics perspective can be represented by the coefficient of performance COP_c ,

$$COP_c = \frac{\dot{P}_{IT}}{\dot{P}_{cool}}. \quad (6.15)$$

Investigators have previously considered the energy-based objective function during joint workload management and cooling control [1, 2, 32, 44]. Our first goal is to maximize COP_c by minimizing \dot{P}_{cool} for a prescribed IT load \dot{P}_{IT} .

6.3.3 Objective 2: Exergy efficiency

In an air-cooled DC equipped with IRC units, the overall exergy loss can be separated into (1) $\dot{\psi}_{ch}$ in the VCR chiller, (2) $\dot{\psi}_{hx}$ in the IRC air-water heat exchangers, (3) $\dot{\psi}_f$ in the fans, and (4) $\dot{\psi}_{serv}$ in the servers. Since the server exergy loss component due to heat generation is determined by the total computing load, which we consider as input data, it cannot be optimized. Therefore, the controllable exergy loss $\dot{\psi}_{cool}$ consists of the three cooling cycle components,

$$\dot{\psi}_{cool} = \dot{\psi}_f + \dot{\psi}_{hx} + \dot{\psi}_{ch}. \quad (6.16)$$

The exergy loss in the chiller is [6, 20],

$$\dot{\psi}_{ch} = \rho_w C_{p,w} \dot{Q}_w T_{amb} \ln\left(\frac{T_{c,w}}{T_{h,w}}\right) + \rho_a C_{p,a} \dot{Q}_{cf} T_{amb} \ln\left(\frac{T_{h,amb}}{T_{amb}}\right). \quad (6.17)$$

Here \dot{Q}_{cf} is the airflow through the condenser fans obtained from the chiller datasheet [37] and $T_{h,amb}$ is the hot air temperature exhausted to the environment through the chiller condenser.

The exergy loss in the IRC heat exchanger,

$$\begin{aligned} \dot{\psi}_{hx} = & \rho_a C_{p,a} \dot{Q}_a \left[(T_{h,a} - T_{c,a}) - T_{amb} \ln \left(\frac{T_{h,a}}{T_{c,a}} \right) \right] \\ & + \rho_w C_{p,w} \dot{Q}_w \left[(T_{c,w} - T_{h,w}) - T_{amb} \ln \left(\frac{T_{c,w}}{T_{h,w}} \right) \right]. \end{aligned} \quad (6.18)$$

That in the fans inside the IRC units is a function of airflow,

$$\dot{\psi}_f = (1 - \eta_f) N_{cu} \left[B_1 + B_2 \left(\frac{\dot{Q}_a}{N_{cu}} \right) + B_3 \left(\frac{\dot{Q}_a}{N_{cu}} \right)^2 \right], \quad (6.19)$$

where the fan efficiency $\eta_f = 0.604$ is obtained from the datasheet of the IRC unit manufactured by RITTAL [36]. The exergy efficiency of the DC cooling system,

$$\eta_{ex} = \left(1 - \frac{\dot{\psi}_{cool}}{\dot{\psi}_{in}} \right) \times 100, \quad (6.20)$$

where the exergy input to the cooling cycle is further evaluated by,

$$\dot{\psi}_{in} = \dot{P}_f + \dot{P}_{ch} + \rho_a C_{p,a} \dot{Q}_a \left[(T_{h,a} - T_{amb}) - T_{amb} \ln \left(\frac{T_{h,a}}{T_{amb}} \right) \right]. \quad (6.21)$$

To reduce thermodynamic irreversibility during operation and maximize effective utilization of cooling capacity, η_{ex} should be maximized [20].

6.3.4 Objective 3: Computing efficiency

The computing efficiency of a server or a processor can be defined in various ways, where we choose the performance per Watt [23, 24], i.e., $\eta_{c,a}$ is calculated as GFLOPs (Giga Floating Point Operations per Second) per unit power consumption in Watts. Although

there are several possible ways of maximizing $\eta_{c,a}$, we focus on the influence of CPU temperature $T_{i,j}^{CPU}$ on $\eta_{c,a}$. Experimental evidence shows that $\eta_{c,a}$ is linearly dependent on $T_{i,j}^{CPU}$, i.e.,

$$\eta_{c,a} = D_1 + D_2 T_{i,j}^{CPU}, \quad (6.22)$$

where the constants $D_1 = 3.0$ and $D_2 = -6.9 \times 10^{-3}$ are obtained from a benchmark case study performed by Druzhinin et al. [24] for Intel Xeon E5-2697 v3 processors.

The maximum achievable computing efficiency $\eta_{c,i}$ of a DC depends on computing cores per CPU, maximum frequency in GHz, FLOPs/cycle of operation and thermal design power (TDP) of the computing processor,

$$\eta_{c,i} = \frac{\text{racks} \times \frac{\text{servers}}{\text{rack}} \times \frac{\text{CPUs}}{\text{server}} \times \frac{\text{cores}}{\text{CPU}} \times \frac{\text{cycles}}{\text{second}} \times \frac{\text{FLOPs}}{\text{cycle}}}{\text{racks} \times \frac{\text{servers}}{\text{rack}} \times \frac{\text{CPUs}}{\text{server}} \times \text{TDP}}, \quad (6.23)$$

A dimensionless metric, computing performance ratio CPR , can be calculated using the ratio of $\eta_{c,a}$ to $\eta_{c,i}$,

$$CPR = \frac{\eta_{c,a}}{\eta_{c,i}}. \quad (6.24)$$

An increase in $T_{i,j}^{CPU}$ causes $\eta_{c,a}$ to decrease which leads to a corresponding decrease in CPR (see Eqs. (6.22) and (6.24)). This reduction can be represented by linear, quadratic, or cubic polynomials that depend on the type and model of the CPU [25]. To deliver a fixed computing performance (in GFLOPs), a lower CPR corresponds to higher computing power (in Watts). Hence, CPR quantifies the loss in computing performance per unit power consumption due to CPU thermal throttling. Other means to optimize CPR include changing the CPU clock speed (GHz) and supply voltage through dynamic voltage and

frequency scaling (DVFS) [25, 26]. The metric CPR is meaningful when the DC has an imposed IT workload, i.e., $0 < LF \leq 1$, where $LF = \frac{\sum_{j=1}^{N_z} \sum_{i=1}^{N_r} u_{i,j}^z}{\sum_{j=1}^{N_z} \sum_{i=1}^{N_r} u_{i,j}^{max}}$. Although $LF = 0$, or zero IT load, provides the highest CPR , that condition is not realistic.

To determine CPR , it is essential to know the mean CPU temperature of each zone $T_{i,j}^{CPU}$ from the thermal model. We utilize a non-linear regression-based auxiliary equation for $T_{i,j}^{CPU}$ [44],

$$T_{i,j}^{CPU} = E_1 + E_2 u_{i,j}^z + E_3 T_{i,j}^f + E_4 (u_{i,j}^z)^2 + E_5 (u_{i,j}^z T_{i,j}^f), \quad (6.25)$$

where $u_{i,j}^z$ and $T_{i,j}^f$ are available from the zonal thermal model. The constants in Eq. (6.25) are determined by performing a non-linear regression using the data obtained from the datasheet [42] for an HP ProLiant DL360 G5 server with two Intel Xeon E5-2697 v3 processors. The resulting values are $E_1 = 20.2$, $E_2 = 59.1$, $E_3 = 0.76$, $E_4 = -8.1$, and $E_5 = -0.56$.

Finally, the mean computing performance ratio CPR_m of the five-rack DC is calculated by taking the average of the $CPRs$ of all the zones,

$$CPR_m = F_1 + F_2 \frac{\sum_{j=1}^{N_z} \sum_{i=1}^{N_r} T_{i,j}^{CPU}}{N_z N_r}, \quad (6.26)$$

where the constants $F_1 = 0.75$, and $F_2 = -1.72 \times 10^{-3}$ are determined by performing regression using the data obtained from [24]. Intuitively, higher CPU temperatures should increase the leakage current of a computing chip and decrease $\eta_{c,a}$ [23], resulting in lower CPR_m . Therefore, the aim is to maximize the value of CPR_m by jointly optimizing workload distribution and cooling system operating parameters.

6.3.5 Optimization problem formulation

Our objectives are to maximize (1) COP_c , (2) η_{ex} , and (3) CPR_m for the five-rack DC with two IRC units. To obtain the optimal cooling parameters and two-dimensional workload distributions, we solve four different optimization problems. The first three optimization problems employ three different objective functions, while in the fourth optimization problem, all three objectives are optimized simultaneously. The structures of the optimization problems are presented in Table 6.3. The essential decision variables that regulate the objective functions are (1) distribution of zonal workload $u_{i,j}^z$, (2) chilled water temperature setpoint $T_{c,w}$, and (3) IRC unit airflow setpoint \dot{Q}_a . Two classes of servers are considered here, (1) active servers with $0 < u_{i,j}^z \leq 1$ and (2) idle servers with $u_{i,j}^z = 0$. Turning OFF an idle server, i.e., making it a passive server to save energy, changes the optimization framework significantly so that the continuous optimization problem becomes a combinatorial optimization problem [1]. Therefore, this consideration is beyond the scope of the present work. The total chilled water flow \dot{Q}_w is fixed to a constant value of 15 GPM, since the optimal values are not sensitive to the water flow [21].

The optimization problems are multi-dimensional with both linear and nonlinear constraints. The first three optimization problems have single objectives and are solved using the *fmincon* function of MATLAB with a *sequential quadratic programming* solver. The fourth optimization problem has three objective functions and is solved using the MOGA-based solver *gamultiobj* in MATLAB. The first constraint for all the optimization problems is linear, ensuring that the DC load factor (LF), i.e., the ratio of current IT load

to maximum capacity $u_{i,j}^{max}$, is set to a particular preset value. The second constraint ensures that the maximum air temperature in the cold (front) chamber is capped below the American society of Heating, Refrigeration, and Air-Conditioning Engineers (ASHRAE) guideline temperature $T_g = 27$ °C [45], which prevents server overheating, ensuring thermally reliable operation. The working range of the decision variables is obtained from the component datasheets of the IRC unit [36] and VCR chiller [37]. For all the optimization problems, the convergence criterion is set to 10^{-6} .

Table 6.3: The four different optimization problems considered.

Optimization 1	Optimization 2
maximize $[COP_c]$ $u_{i,j}^z, T_{c,w}, \dot{Q}_a$	maximize $[\eta_{ex}]$ $u_{i,j}^z, T_{c,w}, \dot{Q}_a$
s.t. $\frac{\sum_{j=1}^{N_z} \sum_{i=1}^{N_r} u_{i,j}^z}{\sum_{j=1}^{N_z} \sum_{i=1}^{N_r} u_{i,j}^{max}} = LF$	s.t. $\frac{\sum_{j=1}^{N_z} \sum_{i=1}^{N_r} u_{i,j}^z}{\sum_{j=1}^{N_z} \sum_{i=1}^{N_r} u_{i,j}^{max}} = LF$
$\max(T_{i,j}^f) \leq T_g$	$\max(T_{i,j}^f) \leq T_g$
$10 \leq T_{c,w}(\text{°C}) \leq 20$	$10 \leq T_{c,w}(\text{°C}) \leq 20$
$1300 \leq \dot{Q}_a(\text{CFM}) \leq 2300$	$1300 \leq \dot{Q}_a(\text{CFM}) \leq 2300$
$0 \leq u_{i,j}^z \leq u_{i,j}^{max}, i = 1, \dots, N_r, \text{ and } j = 1, \dots, N_z$	$0 \leq u_{i,j}^z \leq u_{i,j}^{max}, i = 1, \dots, N_r, \text{ and } j = 1, \dots, N_z$
Optimization 3	Optimization 4
maximize $[CPR_m]$ $u_{i,j}^z, T_{c,w}, \dot{Q}_a$	maximize $[COP_c, \eta_{ex}, CPR_m]$ $u_{i,j}^z, T_{c,w}, \dot{Q}_a$
s.t. $\frac{\sum_{j=1}^{N_z} \sum_{i=1}^{N_r} u_{i,j}^z}{\sum_{j=1}^{N_z} \sum_{i=1}^{N_r} u_{i,j}^{max}} = LF$	s.t. $\frac{\sum_{j=1}^{N_z} \sum_{i=1}^{N_r} u_{i,j}^z}{\sum_{j=1}^{N_z} \sum_{i=1}^{N_r} u_{i,j}^{max}} = LF$
$\max(T_{i,j}^f) \leq T_g$	$\max(T_{i,j}^f) \leq T_g$
$10 \leq T_{c,w}(\text{°C}) \leq 20$	$10 \leq T_{c,w}(\text{°C}) \leq 20$
$1300 \leq \dot{Q}_a(\text{CFM}) \leq 2300$	$1300 \leq \dot{Q}_a(\text{CFM}) \leq 2300$
$0 \leq u_{i,j}^z \leq u_{i,j}^{max}, i = 1, \dots, N_r, \text{ and } j = 1, \dots, N_z$	$0 \leq u_{i,j}^z \leq u_{i,j}^{max}, i = 1, \dots, N_r, \text{ and } j = 1, \dots, N_z$

6.4 Results and discussion

6.4.1 Impact of workload distribution and cooling parameters on front (cold) chamber temperature distribution

The temperature distribution in the cold chamber of an air-cooled DC is an essential feature of the cooling infrastructure. The thermal reliability of a server depends on the cold air intake temperature, which is very sensitive to the zonal utilization levels of servers $u_{i,j}^z$ and cooling parameters $T_{c,w}$ and \dot{Q}_a [21]. Instead of examining the entire cold chamber temperature to assess thermal characteristics, we consider the metric active server temperature distribution (*ASTD*) [38],

$$ASTD = \text{mean}(T_{i,j}^f) + \text{stdev}(T_{i,j}^f). \quad (6.27)$$

The cold air temperature $T_{i,j}^f$ supplied by the IRC units depends on both $T_{c,w}$ and \dot{Q}_a for a fixed \dot{Q}_w , whereas the spatial variance of $T_{i,j}^f$ is regulated by \dot{Q}_a only. Therefore, we note that *ASTD* can be regulated by altering $T_{c,w}$, \dot{Q}_a , and the distribution of $u_{i,j}^z$.

Table 6.4 and Figure 6.3 show the influence of jointly varying $T_{c,w}$, \dot{Q}_a , and the distribution of $u_{i,j}^z$ on the *ASTD* and rack-height averaged cold chamber temperature distributions. Comparing cases A, B, and C reveals that simultaneously increasing $T_{c,w}$ and \dot{Q}_a raises *ASTD*, bringing it closer to the maximum allowable cold chamber temperature of 27 °C [45]. This, in turn, minimizes the overcooling of servers in the vicinity of IRC units, as shown in Figure 6.3 (Cases A and C). Comparing cases C, D, and E shows that workload distribution has an insignificant effect on the *ASTD* and temperature distribution.

Table 6.4: Parametric details of the cases presented in Figure 6.3 with $LF = 0.5$ and $\max(T_{i,j}^f) \leq 27^\circ\text{C}$. Red and dark green colors are mapped for two extreme values of $u_{i,j}^z$, 1 and 0, respectively.

Case	Distribution of $u_{i,j}^z$ across zones	$T_{c,w}$ ($^\circ\text{C}$)	\dot{Q}_a (CFM)	$ASTD$ ($^\circ\text{C}$)						
A	Uniform workload assignment					20	2153	26.04		
	IRC unit (Left)	R1	R2	R3	R4				R5	IRC unit (Right)
		0.5	0.5	0.5	0.5				0.5	
		0.5	0.5	0.5	0.5				0.5	
		0.5	0.5	0.5	0.5				0.5	
		0.5	0.5	0.5	0.5				0.5	
B	Uniform workload assignment					10	1754	23.08		
	IRC unit (Left)	R1	R2	R3	R4				R5	IRC unit (Right)
		0.5	0.5	0.5	0.5				0.5	
		0.5	0.5	0.5	0.5				0.5	
		0.5	0.5	0.5	0.5				0.5	
		0.5	0.5	0.5	0.5				0.5	
C	Uniform workload assignment					15	1930	24.53		
	IRC unit (Left)	R1	R2	R3	R4				R5	IRC unit (Right)
		0.5	0.5	0.5	0.5				0.5	
		0.5	0.5	0.5	0.5				0.5	
		0.5	0.5	0.5	0.5				0.5	
		0.5	0.5	0.5	0.5				0.5	
D	Workload assigned away from IRC units					15	1978	24.36		
	IRC unit (Left)	R1	R2	R3	R4				R5	IRC unit (Right)
		0	0.75	1	0.75				0	
		0	0.75	1	0.75				0	
		0	0.75	1	0.75				0	
		0	0.75	1	0.75				0	
E	Workload assigned close to IRC units					15	1867	24.77		
	IRC unit (Left)	R1	R2	R3	R4				R5	IRC unit (Right)
		1	0.25	0	0.25				1	
		1	0.25	0	0.25				1	
		1	0.25	0	0.25				1	
		1	0.25	0	0.25				1	

To identify local hotspots in the multi-rack DC, we provide rack-height averaged temperature distributions for the cases presented in Table 6.4. The red dotted line in Figure 6.3 corresponds to the maximum allowable temperatures in each rack [45]. Since the 3rd rack (R3) is located farthest from both the IRC units, it experiences the highest cold air temperature. Two physical phenomena are identified as primary causes for the hotspot in rack R3: (1) this hotspot has a higher airflow path resistance being farthest from the IRC units (see Figure 6.2), which increases the temperature, and (2) the lower air supply from the IRC units towards R3 increases the possibility of hot-air recirculation from the back (hot) chamber to the front (cold) chamber, thus increasing the local temperature. Overall, the cold chamber temperature is an essential feature of the DC optimization framework since the maximum cold chamber temperature influences thermal reliability.

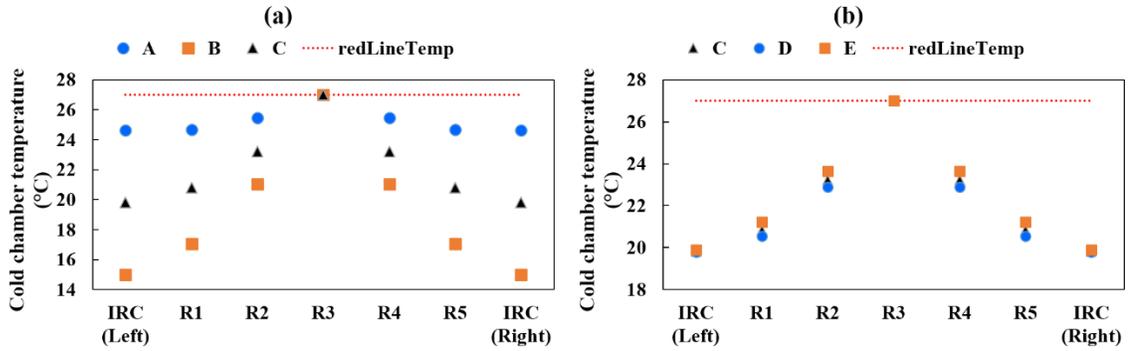


Figure 6.3: Influence of (a) cooling parameters $T_{c,w}$ and \dot{Q}_a and (b) workload assignment on the rack height averaged cold chamber temperature distribution for the different cases listed in Table 6.4. The red dotted line in the figures denotes the maximum allowable temperature in the cold chamber.

6.4.2 Trade-offs in energy, exergy, and computing efficiency

To account for the interplay among different decision variables and objective functions, we first examine the parametric trade-offs across essential quantities such as (1) $u_{i,j}^z$, (2)

$T_{c,w}$, (3) \dot{Q}_a , (4) LF , (5) COP_c , (6) η_{ex} , and (7) CPR_m , presented in Tables 6.5 and 6.6. For all cases, the maximum cold air temperature at the server inlets is maintained below 27 °C and $\dot{Q}_w = 15$ GPM. Table 6.5 shows the influence of coordinated regulation of $u_{i,j}^z$, $T_{c,w}$, and \dot{Q}_a for $LF = 0.5$. For a constant LF , raising $T_{c,w}$ from 10 °C to 20 °C requires approximately 26% higher \dot{Q}_a to maintain the maximum server inlet temperature below 27 °C. Doing so reduces the spatial air temperature variances in the cold chamber and improves η_{ex} by up to 3.5%. Although raising $T_{c,w}$ requires larger \dot{Q}_a , the overall COP_c of the DC improves by approximately 14% when $T_{c,w}$ is increased from 10 °C to 20 °C. This can be attributed to the fact that the energy savings obtained by operating the VCR chiller at a higher $T_{c,w}$ is significantly larger than what is required to drive the IRC fans [36, 37].

Assigning the server workloads closer to the IRC units (racks R1 and R5) requires up to 4.5% less \dot{Q}_a due to the corresponding reduction in flow path resistance, which slightly improves COP_c (by approximately 1.5%), although η_{ex} is virtually unaltered when the distribution of $u_{i,j}^z$ changes. The mean computing performance ratio CPR_m is not significantly affected by jointly regulating the cooling parameters and workload distribution for a specific LF . Since the server inlet air temperature $T_{i,j}^f$ weakly influences $T_{i,j}^{CPU}$, CPR_m is unaltered [41]. Therefore, only two objectives, COP_c and η_{ex} , are controlled by altering the three decision variables $u_{i,j}^z$, $T_{c,w}$, and \dot{Q}_a .

Table 6.5: Influence of zonal workload assignments and cooling cycle decision variables on the three objectives for $LF = 0.5$ while satisfying the constraint $\max(T_{i,j}^f) \leq 27$ °C. For all the cases $\dot{Q}_w = 15$ GPM. Red and dark green colors are mapped for the extreme values of $u_{i,j}^z$, 1 and 0, respectively.

Distribution of $u_{i,j}^z$ across zones							$T_{c,w}$ (°C)	\dot{Q}_a (CFM)	COP_c	η_{ex} (%)	CPR_m
Case1: Uniformly distributed workload							10	1754	3.62	11.0	0.650
							12	1820	3.74	11.1	0.650
IRC unit (Left)	R1	R2	R3	R4	R5	IRC unit (Right)	14	1891	3.85	11.5	0.649
	0.5	0.5	0.5	0.5	0.5		16	1970	3.95	12.0	0.648
	0.5	0.5	0.5	0.5	0.5		18	2057	4.05	12.9	0.647
	0.5	0.5	0.5	0.5	0.5		20	2153	4.13	14.2	0.646
	0.5	0.5	0.5	0.5	0.5						
	0.5	0.5	0.5	0.5	0.5						
Case2: Workload assigned close to IRC units							10	1674	3.66	10.1	0.651
							12	1745	3.78	10.3	0.651
IRC unit (Left)	R1	R2	R3	R4	R5	IRC unit (Right)	14	1824	3.89	10.7	0.650
	1	0.25	0	0.25	1		16	1912	4.00	11.4	0.650
	1	0.25	0	0.25	1		18	2010	4.09	12.4	0.649
	1	0.25	0	0.25	1		20	2121	4.17	13.9	0.649
	1	0.25	0	0.25	1						
	1	0.25	0	0.25	1						
Case 3: Workload assigned away from IRC units							10	1816	3.59	11.7	0.655
							12	1877	3.70	11.8	0.654
IRC unit (Left)	R1	R2	R3	R4	R5	IRC unit (Right)	14	1943	3.81	12.1	0.652
	0	0.75	1	0.75	0		16	2014	3.92	12.6	0.652
	0	0.75	1	0.75	0		18	2092	4.02	13.4	0.650
	0	0.75	1	0.75	0		20	2179	4.11	14.5	0.650
	0	0.75	1	0.75	0						
	0	0.75	1	0.75	0						

The influence of LF on the different objectives and decision variables is examined in Table 6.6. For all cases, the maximum temperature in the front chamber is maintained below 27 °C and \dot{Q}_w is maintained at 15 GPM. Increasing LF from 0.2 to 0.8 shows that COP_c and η_{ex} improve by 23% and 6%, respectively, while the CPR_m degrades by

approximately 6%. The degradation in the computing performance ratio CPR_m can be attributed to the elevated CPU temperature at higher LF s. As the VCR chiller operates efficiently from both energy and exergy perspectives at high LF s, both the 1st and 2nd law efficiencies improve, which improves overall COP_c and η_{ex} . Since LF is not under the control of the DC operator, CPR_m cannot be regulated by homogenous workload management across DC racks. Therefore, heterogeneous LF management across multi-rack DC systems should be adopted.

Tables 6.5 and 6.6 show that the three objectives are predominantly functions of $T_{c,w}$, \dot{Q}_a , and the distribution of $u_{i,j}^z$. Therefore, instead of adopting a brute force parametric variation approach, in subsequent sections, we solve non-linear single- and multi-objective optimization problems to determine the optimum decision variables.

Table 6.6: Influence of LF on decision variables and objectives while satisfying the constraint $\max(T_{i,j}^f) \leq 27$ °C. For all cases $\dot{Q}_w = 15$ GPM.

Distribution of $u_{i,j}^z$ across zones							$T_{c,w}$ (°C)	\dot{Q}_a (CFM)	COP_c	η_{ex} (%)	CPR_m
Case 1: $LF = 0.2$							10	1661	3.42	10.2	0.673
							12	1730	3.53	10.3	0.672
IRC unit (Left)	R1	R2	R3	R4	R5	IRC unit (Right)	14	1807	3.63	10.7	0.671
	0.2	0.2	0.2	0.2	0.2		16	1891	3.72	11.3	0.669
	0.2	0.2	0.2	0.2	0.2		18	1985	3.80	12.3	0.668
	0.2	0.2	0.2	0.2	0.2		20	2091	3.86	13.7	0.667
	0.2	0.2	0.2	0.2	0.2		10	1754	3.62	11.0	0.651
	0.2	0.2	0.2	0.2	0.2		12	1820	3.74	11.1	0.650
Case 2: $LF = 0.5$							14	1891	3.85	11.5	0.649
							16	1970	3.96	12.1	0.648
IRC unit (Left)	R1	R2	R3	R4	R5	IRC unit (Right)	18	2057	4.05	12.9	0.647
	0.5	0.5	0.5	0.5	0.5		20	2153	4.14	14.2	0.647
	0.5	0.5	0.5	0.5	0.5		10	1829	3.75	11.5	0.632
	0.5	0.5	0.5	0.5	0.5		12	1891	3.88	11.6	0.631
	0.5	0.5	0.5	0.5	0.5		14	1959	4.00	12.0	0.631
	0.5	0.5	0.5	0.5	0.5		16	2033	4.12	12.5	0.630
Case 3: $LF = 0.8$							18	2113	4.23	13.4	0.629
							20	2211	4.32	14.6	0.629
IRC unit (Left)	R1	R2	R3	R4	R5	IRC unit (Right)	10	1829	3.75	11.5	0.632
	0.8	0.8	0.8	0.8	0.8		12	1891	3.88	11.6	0.631
	0.8	0.8	0.8	0.8	0.8		14	1959	4.00	12.0	0.631
	0.8	0.8	0.8	0.8	0.8		16	2033	4.12	12.5	0.630
	0.8	0.8	0.8	0.8	0.8		18	2113	4.23	13.4	0.629
	0.8	0.8	0.8	0.8	0.8		20	2211	4.32	14.6	0.629

6.4.3 Single-objective optimization

6.4.3.1 Maximization of COP_c

The energy-based metric COP_c can be optimized by altering the cooling parameters, $T_{c,w}$ and \dot{Q}_a , and the zonal workload distribution $u_{i,j}^z$. The explicit optimization problem is

shown in Table 6.3. The total chilled water flowrate \dot{Q}_w is fixed at 15 GPM. The aim of the COP_c maximization problem is to find optimal values of $T_{c,w}$, \dot{Q}_a , and $u_{i,j}^z$ while maintaining the cold chamber temperature below 27 °C. Optimum values of decision variables and the objective are presented in Table 6.7 for five different LF s, 0.2, 0.35, 0.5, 0.65, and 0.8.

The COP_c maximization problem assigns the server workload closest to the IRC units for the entire range of LF values [1, 21]. As LF increases, the servers located near IRC units (in racks R1 and R5) are fully utilized, i.e., $u_{i,j}^z = 1$, which necessitates assigning additional incoming workloads to racks away from the IRC units (R2, R4, and R3). This finding is in line with prior one-dimensional workload management algorithms that minimize the cooling energy consumption in air-cooled DCs [1]. However, we reveal the importance of the heterogeneity of $u_{i,j}^z$ inside a rack that creates a two-dimensional workload variation, which is essential for hotspot management. COP_c maximization results in a relatively high $T_{c,w} = 20$ °C, which is invariant with changing LF . Raising $T_{c,w}$ reduces the energy consumption component of the chiller and maximizes COP_c . However, the total airflow \dot{Q}_a of the fans inside IRC units vary non-linearly (see Eq. 6.13) with changing LF . A higher LF results in higher heat generation from servers, when \dot{Q}_a must be raised to bring the maximum cold chamber temperature below 27 °C. A higher LF significantly improves COP_c (by up to 12.5%), since the VCR chiller operates more effectively at higher LF s [33]. Although COP_c maximization reduces cooling energy consumption; it does not necessarily diminish thermodynamic irreversibility in the DC cooling system during operation. This often leads to loss of available cooling capacity and further reduces η_{ex} [6, 20, 21].

Table 6.7: Optimal decision variables obtained from the COP_c maximization problem for different LF . Red and dark green colors are mapped for the extreme values of $u_{i,j}^z$, 1 and 0, respectively.

LF	Distribution of $u_{i,j}^z$ across zones						$T_{c,w}$ (°C)	\dot{Q}_a (CFM)	COP_c	
0.2	IRC unit (Left)	R1	R2	R3	R4	R5	IRC unit (Right)	20	2034	3.92
		0.55	0.00	0.00	0.00	0.55				
		0.41	0.00	0.00	0.00	0.34				
		0.68	0.00	0.00	0.00	0.42				
		0.55	0.00	0.00	0.00	0.48				
0.55	0.00	0.00	0.00	0.46						
0.3 5	IRC unit (Left)	R1	R2	R3	R4	R5	IRC unit (Right)	20	2054	4.08
		0.97	0.00	0.00	0.00	0.77				
		1.00	0.00	0.00	0.00	0.00				
		0.76	0.00	0.00	0.00	1.00				
		0.88	0.00	0.00	0.00	0.88				
0.88	0.00	0.00	0.00	0.86						
0.5	IRC unit (Left)	R1	R2	R3	R4	R5	IRC unit (Right)	20	2081	4.21
		1.00	0.24	0.00	0.21	1.00				
		1.00	0.29	0.00	0.30	1.00				
		1.00	0.30	0.00	0.30	1.00				
		1.00	0.13	0.00	0.14	1.00				
1.00	0.29	0.00	0.30	1.00						
0.6 5	IRC unit (Left)	R1	R2	R3	R4	R5	IRC unit (Right)	20	2107	4.32
		1.00	0.58	0.00	0.77	1.00				
		1.00	0.68	0.00	0.66	1.00				
		1.00	0.34	0.00	1.00	1.00				
		1.00	0.66	0.00	0.69	1.00				
1.00	0.38	0.00	0.50	1.00						
0.8	IRC unit (Left)	R1	R2	R3	R4	R5	IRC unit (Right)	20	2130	4.41
		1.00	1.00	0.00	1.00	1.00				
		1.00	1.00	0.00	1.00	1.00				
		1.00	1.00	0.00	1.00	1.00				
		1.00	1.00	0.00	1.00	1.00				

6.4.3.2 Maximization of η_{ex}

To address the drawback of COP_c maximization from the 2nd law of thermodynamics perspective, the second optimization problem (see Table 6.3) maximizes η_{ex} . The nonlinear constrained optimization is solved using the methodology in Section 2.5. The resulting values of decision variables and optimized objective function values are presented in Table 6.8 for a wide range of LF s. In contrast to the conventional approach of assigning workloads to racks closest to the IRC units, the η_{ex} maximization problem partially distributes the server workload across the racks (see workload color maps in Table 6.8). A high value of \dot{Q}_a is preferred since raising the airflow diminishes $\dot{\psi}_{cool}$ due to spatial air temperature gradients. Increasing \dot{Q}_a also reduces the probability of hot air recirculation and hotspot formation in the cold chamber [21]. A low value of $T_{c,w}$ is preferred since, at lower chilled water temperatures, the 2nd law efficiency of the VCR chiller improves, and the heat transfer efficiency (or effectiveness) of the IRC heat exchanger is enhanced. As the LF changes, the cooling cycle variables \dot{Q}_a and $T_{c,w}$ remain invariant while the workload distribution changes, resulting in η_{ex} maximization. The LF does not significantly affect (less than 2% difference) η_{ex} and the highest value of $\eta_{ex} = 18.6\%$ is observed for an $LF = 0.2$.

Table 6.8: Optimal decision variables obtained from the η_{ex} maximization problem for different LF . Red and dark green colors are mapped for the extreme values of $u_{i,j}^z$, 1 and 0, respectively.

LF	Distribution of $u_{i,j}^z$ across zones						$T_{c,w}$ (°C)	\dot{Q}_a (CFM)	η_{ex} (%)	
0.2	IRC unit (Left)	R1	R2	R3	R4	R5	IRC unit (Right)	10	2300	18. 6
		0.00	0.15	0.03	0.15	0.11				
		0.86	0.15	0.03	0.15	1.00				
		0.01	0.27	0.02	0.27	0.28				
		0.27	0.27	0.02	0.27	0.01				
		0.05	0.27	0.02	0.27	0.05				
0.3 5	IRC unit (Left)	R1	R2	R3	R4	R5	IRC unit (Right)	10	2300	18. 1
		0.15	0.51	0.06	0.55	0.96				
		0.19	0.59	0.00	0.56	0.17				
		0.19	0.25	0.00	0.42	0.97				
		0.21	0.41	0.03	0.21	0.99				
		0.22	0.43	0.05	0.41	0.22				
0.5	IRC unit (Left)	R1	R2	R3	R4	R5	IRC unit (Right)	10	2300	17. 6
		0.66	0.25	0.00	0.55	0.73				
		0.70	0.83	0.00	0.49	0.64				
		0.60	1.00	0.00	1.00	0.68				
		0.46	0.11	0.00	0.12	0.60				
		0.64	0.92	0.00	0.92	0.61				
0.6 5	IRC unit (Left)	R1	R2	R3	R4	R5	IRC unit (Right)	10	2300	17. 1
		0.89	0.41	0.08	0.00	0.70				
		0.69	0.71	0.03	0.01	0.75				
		1.00	1.00	0.98	1.00	1.00				
		1.00	1.00	0.69	1.00	1.00				
		0.26	0.91	0.00	0.93	0.21				
0.8	IRC unit (Left)	R1	R2	R3	R4	R5	IRC unit (Right)	10	2300	16. 7
		1.00	1.00	0.38	1.00	1.00				
		1.00	1.00	0.07	1.00	1.00				
		1.00	1.00	0.00	0.98	1.00				
		1.00	0.00	0.48	0.15	1.00				
		1.00	1.00	0.94	1.00	1.00				

The first two optimization problems maximize COP_c and η_{ex} of the cooling cycle from the perspective of the 1st and 2nd laws of thermodynamics, respectively. These two

approaches reduce the cooling cycle operational expenditure and the loss of effective cooling capacity. However, they do not necessarily reduce overheating of the CPU, which degrades the value of the computing performance ratio CPR_m (see Eqs. (6.22)-(6.26)).

6.4.3.3 Maximization of CPR_m

The third optimization problem maximizes the computing performance per unit energy consumption by maximizing CPR_m (see Table 6.3). It is worth noting that CPR_m only captures the effects of computing performance degradation at elevated chip temperature. This approach is somewhat similar to CPU temperature-aware workload scheduling and cooling control, where the idea is to keep the CPU die temperature below a certain value [11]. The objective function CPR_m is represented as a function of the mean CPU temperatures of each zone $T_{i,j}^{CPU}$ (see Eqs. (6.22)-(6.26)). From Tables 6.5 and 6.6, we note that CPR_m is altered by changing $T_{c,w}$, \dot{Q}_a , the distribution of $u_{i,j}^z$, and LF . Therefore, the aim is to find $T_{c,w}$, \dot{Q}_a , and $u_{i,j}^z$ for a wide range of LF while keeping the cold chamber air temperature below 27 °C. Table 6.9 shows the optimal values of the decision variables and the objective function for various LF scenarios.

Table 6.9: Optimal decision variables obtained from the CPR_m maximization problem for different LF . Red and dark green colors are mapped for the extreme values of $u_{i,j}^z$, 1 and 0, respectively.

LF	Distribution of $u_{i,j}^z$ across zones						$T_{c,w}$ (°C)	\dot{Q}_a (CFM)	CPR_m	
0.2	IRC unit (Left)	R1	R2	R3	R4	R5	IRC unit (Right)	10	2300	0.682
		1.00	0.00	1.00	0.00	0.00				
		1.00	0.00	1.00	0.00	0.00				
		0.00	0.00	1.00	0.00	0.00				
		0.00	0.00	0.00	0.00	0.00				
0.3 5	IRC unit (Left)	R1	R2	R3	R4	R5	IRC unit (Right)	10	2300	0.671
		1.00	0.00	0.00	0.00	0.00				
		1.00	0.00	0.00	0.00	0.75				
		1.00	1.00	0.00	1.00	1.00				
		1.00	0.00	0.00	0.00	1.00				
0.5	IRC unit (Left)	R1	R2	R3	R4	R5	IRC unit (Right)	10	2300	0.659
		1.00	0.00	0.00	1.00	1.00				
		1.00	1.00	0.00	1.00	1.00				
		0.50	0.00	0.00	0.00	1.00				
		1.00	1.00	0.00	1.00	1.00				
0.6 5	IRC unit (Left)	R1	R2	R3	R4	R5	IRC unit (Right)	10	2300	0.648
		1.00	1.00	0.00	1.00	1.00				
		1.00	1.00	0.00	1.00	1.00				
		0.00	0.00	0.25	0.00	0.00				
		1.00	1.00	0.00	1.00	1.00				
0.8	IRC unit (Left)	R1	R2	R3	R4	R5	IRC unit (Right)	10	2300	0.637
		1.00	0.00	1.00	0.00	1.00				
		1.00	1.00	1.00	0.00	1.00				
		1.00	1.00	0.00	1.00	1.00				
		1.00	1.00	1.00	1.00	1.00				

A low $T_{c,w}$ and high \dot{Q}_a are preferred since regulating these two variables reduces the cold air supply temperature to the servers $T_{i,j}^f$. As a result of lower cold air temperature

at the server inlet, the CPU temperature decreases and improves CPR_m by less than 1%. The approach also distributes workload partially across the racks rather than assigning them closest to the IRC units for a wide range of LF s, as evident from the two-dimensional heatmaps of zonal utilization distribution (see Table 6.9). For the majority of zones, either $u_{i,j}^z = 0$ or $u_{i,j}^z = 1$ is preferred. Raising LF degrades the maximum achievable CPR_m since a larger LF is associated with elevated CPU temperatures (> 65 °C). For example, increasing LF from 0.2 to 0.8 diminishes the CPR_m by approximately 6.6% as CPU temperatures increase with higher LF s. Thus, CPR_m is a weak function of $T_{c,w}$ and \dot{Q}_a , while it is strongly influenced by LF and the distribution of $u_{i,j}^z$ [41].

6.4.3.4 Summary of single-objective optimization results

The results obtained by solving the COP_c , η_{ex} , and CPR_m single-objective maximization problems for different values of LF are presented in Figure 6.4. As the LF for the DC increases, COP_c improvement is possible due to the higher energy efficiency of the VCR chiller [37]. An elevated LF raises the cooling load on the VCR chiller, moving it towards the designed cooling capacity of the chiller. On the other hand, changing LF has an insignificant effect (less than a 2% difference) on the maximum achievable η_{ex} , i.e., the exergy efficiency from the single-objective optimization problem is essentially independent of LF . The third optimization problem, CPR_m maximization, shows the influence of LF on the mean computing performance (see Figure 6.4a). Increasing LF from 0.2 to 0.8 degrades the maximum achievable CPR_m by approximately 7.5% as $T_{i,j}^{CPU}$ is elevated. The lower values of CPR_m correspond to poor computing performance, i.e., lower

GFLOPs per unit power consumption, resulting in higher latencies and increased energy expenditure. The result is that the three different objectives cannot be achieved with similar combinations of decision variables and LF s. To address this problem, we solve a non-linear tri-objective optimization problem to address the trade-offs across objective functions and decision variables using a MOGA-based Pareto front approach.

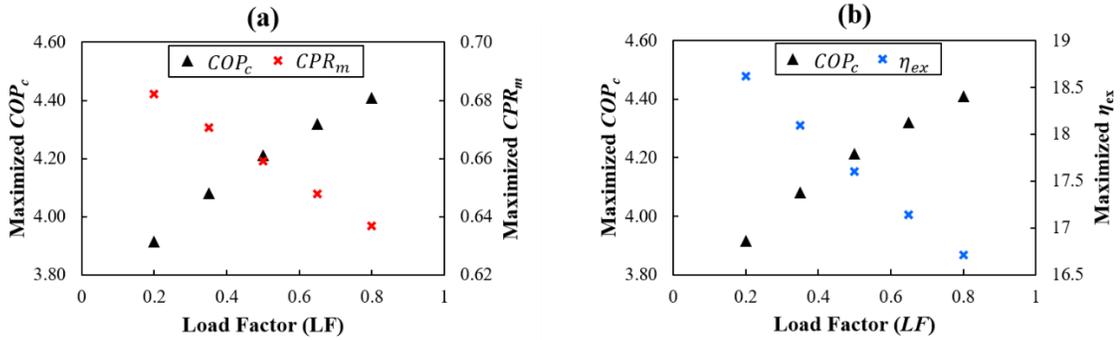


Figure 6.4: Effect of load factor LF on optimum values of three objective functions: (a) Variation of COP_c and CPR_m as a function of LF , (b) Variation of COP_c and η_{ex} as a function of LF .

6.4.4 Multi-objective optimization

To address the trade-offs between the three different single-objective optimization problems, the fourth optimization problem is MOGA-based, incorporating knowledge of all the three objective functions simultaneously (see Table 6.3). The tri-objective optimization problem is solved for five different LF s, 0.2, 0.35, 0.5, 0.65, and 0.8 leading to a three-dimensional Pareto front (or trade-off diagram). To better represent the relation between different objectives, we separate the three-dimensional Pareto front in Figure 6.5 into three different two-dimensional projections showing trends of (1) η_{ex} vs. COP_c , (2) CPR_m vs. COP_c , and (3) CPR_m vs. η_{ex} . In all trade-off diagrams, the “desired region” and the “undesired region” correspond to the regimes where both the objectives are

simultaneously maximized and minimized, respectively. For each LF , a separate Pareto front is obtained, as shown in Figures 6.5a, b, and c. Figure 6.5a reveals a trade-off between the energy-based metric COP_c and exergy-based metric η_{ex} . High values for both COP_c and η_{ex} can only be achieved with higher LF when the VCR chiller operates close to its designed cooling capacity. As LF is increased from 0.2 to 0.8, the Pareto front shifts towards the desired region. Figure 6.5b suggests that when the LF is increased from 0.2 to 0.8, even though the COP_c improves by 20%, the CPR_m lowers by approximately 7.5%. Operation at a lower LF (for example, $LF = 0.2$) stretches the Pareto optimal front of COP_c so that the DC operator can make the desired choice across a wider range than for a higher LF . Figure 6.5c shows the proportionality trade-off between CPR_m and η_{ex} . At a certain LF , a high value of η_{ex} (approximately 18%) is achievable by regulating the cooling cycle parameters $T_{c,w}$ and \dot{Q}_a , which slightly enhances the CPR_m (up to 1%).

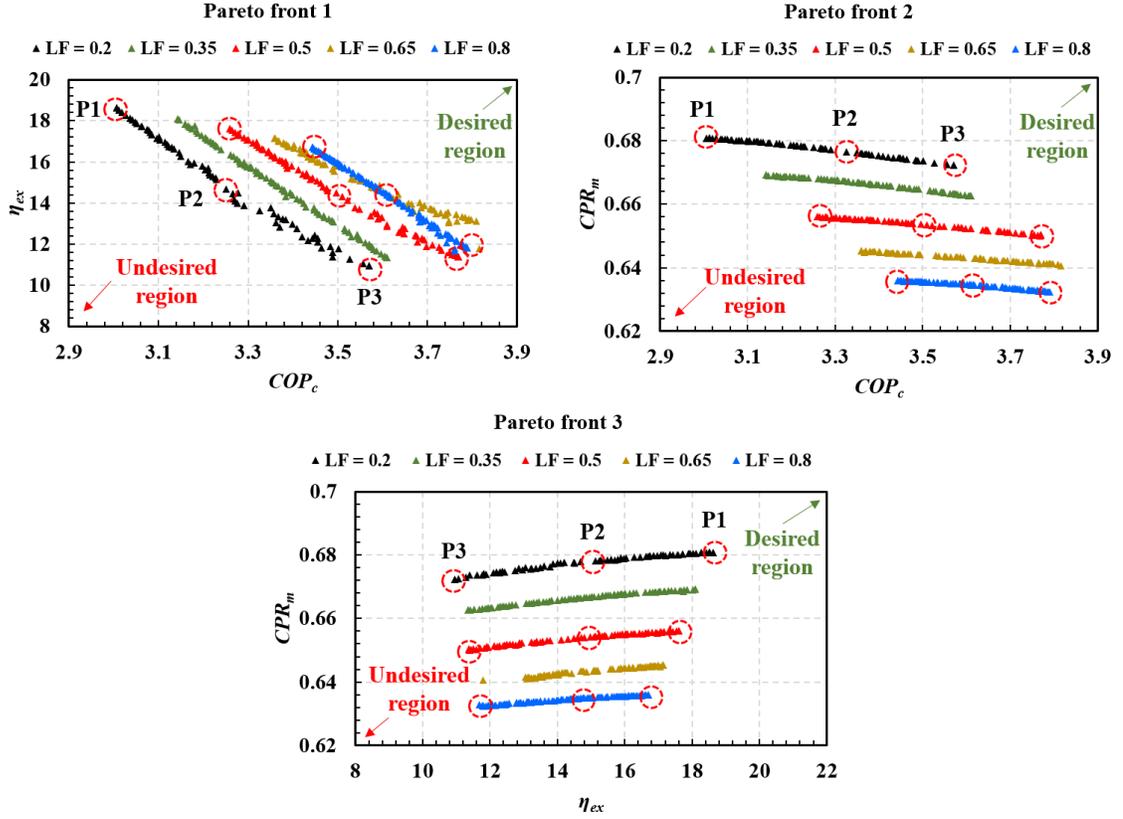


Figure 6.5: Pareto optimal fronts obtained by solving the multi-objective optimization problem while satisfying $T_g = 27^\circ\text{C}$ for different LFs. The tri-axis Pareto front is separated into three different two-dimensional Pareto fronts: (1) η_{ex} vs. COP_c , (2) CPR_m vs. COP_c , and (3) CPR_m vs. η_{ex} for better representing the trade-offs across different objectives. The “desired region” and the “undesired region” are the regimes where both the objectives are simultaneously maximized and minimized, respectively.

To gain better insight into the interplay across decision variables and thermal behavior, we present the values of $T_{c,w}$, \dot{Q}_a , CPR_m , COP_c , η_{ex} , distributions of $u_{i,j}^z$, $ASTD$ and rack-height averaged temperature distribution for salient points across the Pareto fronts in Table 6.10, and Figures 6.6 and 6.7.

Table 6.10: Decision variables and objective function values for nine salient points marked on the Pareto optimal fronts presented in Figure 6.5.

LF	Point	$T_{c,w}$ (°C)	\dot{Q}_a (CFM)	COP_c	η_{ex} (%)	CPR_m	$ASTD$ (°C)
0.2	P1	10.0	2300	3.0	18.6	0.680	13.9
	P2	10.7	2056	3.2	15.1	0.678	16.9
	P3	13.2	1804	3.6	10.9	0.672	23.2
0.5	P1	10.0	2300	3.2	17.6	0.656	14.7
	P2	11.0	2074	3.5	14.5	0.653	18.2
	P3	12.7	1855	3.8	11.4	0.650	23.6
0.8	P1	10.0	2300	3.4	16.7	0.635	15.4
	P2	10.4	2092	3.6	14.3	0.634	18.3
	P3	10.8	1877	3.8	11.8	0.632	22.8

Moving from P1 towards P3 on a particular Pareto front slightly affects the distribution of $u_{i,j}^z$, where P3 is associated with more server zones at $u_{i,j}^z = 1$. However, changing LF entirely alters the workload distribution to support a higher IT load (see Figure 6.6). Nevertheless, Table 6.10 shows that the cooling parameters $T_{c,w}$ and \dot{Q}_a regulate the desired values of COP_c and η_{ex} for a specified LF . Moving from P1 towards P3 for a specific LF results in an increase in $T_{c,w}$ while simultaneously decreasing \dot{Q}_a . This raises the mean and variance of the cold chamber air temperature, as evident from the $ASTD$ and the rack height average temperature distribution presented in Table 6.10 and Figure 6.7, respectively. Although raising the mean cold air temperature is beneficial from the perspective of maximizing COP_c , increasing the spatial variance in the air temperature causes local hotspots, leading to a higher system irreversibility and reduced η_{ex} . Therefore, an intermediate point on the Pareto front can yield a desirable trade-off across COP_c and η_{ex} . As the LF increases, the extent of variation for $T_{c,w}$ and \dot{Q}_a decreases, consequently yielding higher values of COP_c and η_{ex} . Although $ASTD$ is altered upon moving from P1

to P3 on a Pareto front, it is nearly unaffected by changing LF . Therefore, we conclude that LF has an insignificant effect on the cold chamber temperature distribution (see Figure 6.7). The third objective, CPR_m , changes only by 1% on a Pareto front for a given LF . However, increasing the LF degrades the value of CPR_m by as much as 7%.

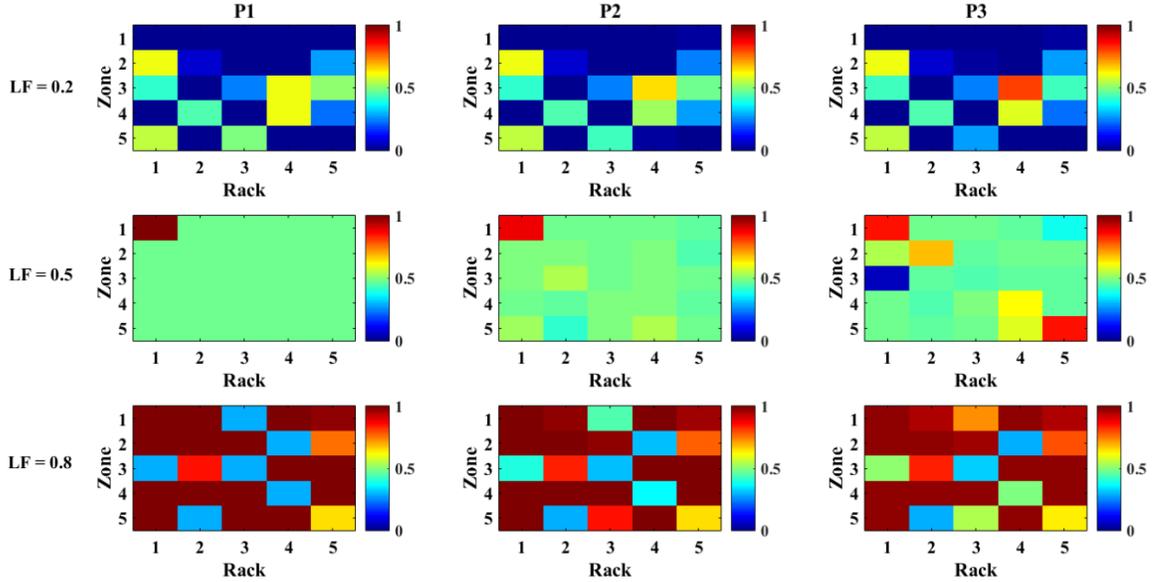


Figure 6.6: Two-dimensional distributions of zonal utilizations $u_{i,j}^z$ for nine different points marked on the Pareto fronts shown in Figure 6.5. The color bars are scaled for the extreme values of $u_{i,j}^z$, 0 and 1, respectively. This figure is to be interpreted in accordance with Table 6.10.

Examining the rack-height averaged temperature distribution (see Figure 6.7) for salient points on the Pareto fronts shows that for all LF s considered, the regime where P1 lies is associated with low cold chamber air temperature and lower variance of temperature across the racks. On the other hand, the P3 points for different LF curves show a high spatial variance of the cold air temperature. The maximum temperature is observed in rack R3 since it is located at the farthest distance from both the IRC units, resulting in a higher flow path resistance, as shown by the FNR (see Figure 6.2). Therefore, the largest pressure

drop is experienced by this rack, leading to an increased probability of hot air recirculation from the hot chamber to the cold chamber via the air-blocking brushes.

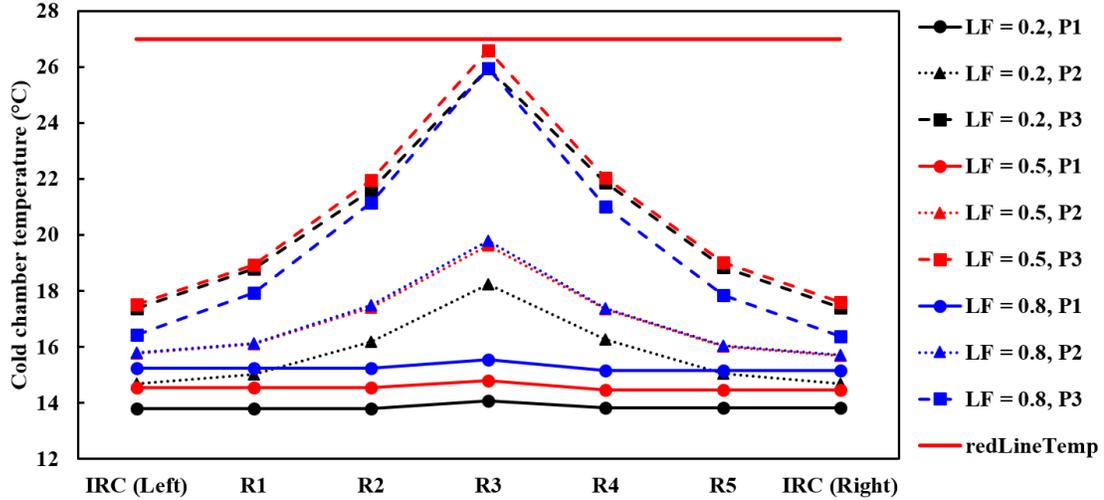


Figure 6.7: Rack height averaged cold chamber temperature distributions for the nine salient points marked on the Pareto optimal fronts presented in Figure 6.5. The “redLineTemp” corresponds to the 27 °C maximum allowable cold air temperature for servers. This figure is to be interpreted in accordance with Table 6.10 and Figure 6.6.

6.4.5 CPR improvement strategy for large-scale DCs

Upon examining the Pareto fronts presented in Figure 6.5, it is evident that desired values of η_{ex} and COP_c are obtained by appropriately selecting the cooling parameters $T_{c,w}$ and \dot{Q}_a , reducing operational energy expenditure by maximizing COP_c and enhancing effective cooling capacity by maximizing η_{ex} . However, at a higher DC workload (i.e., larger LF), there is up to 7.5% degradation in CPR_m due to elevated CPU temperatures. Using an exhaustive search, we found that this degradation could not be avoided. Therefore, further effort is required to improve CPR_m for large-scale DCs which, in turn, will improve $\eta_{c,a}$. Fitting CPR_m with respect to LF (Figure 6.4a), results in the following relation,

$$CPR_m = -0.075(LF) + 0.697, \quad (6.28)$$

which is valid for the range $0 < LF \leq 1$.

Next, we assume a large-scale DC infrastructure with a total number of IT racks $N_{total} = 100$ that can be constructed by several identical modular blocks of enclosures or “DC pods.” Each pod consists of (1) five air-cooled DC racks with a standalone enclosure, (2) two IRC units, (3) a water circulation pump, (4) a VCR chiller, and (5) a pod workload/cooling manager (see Figure 6.8).

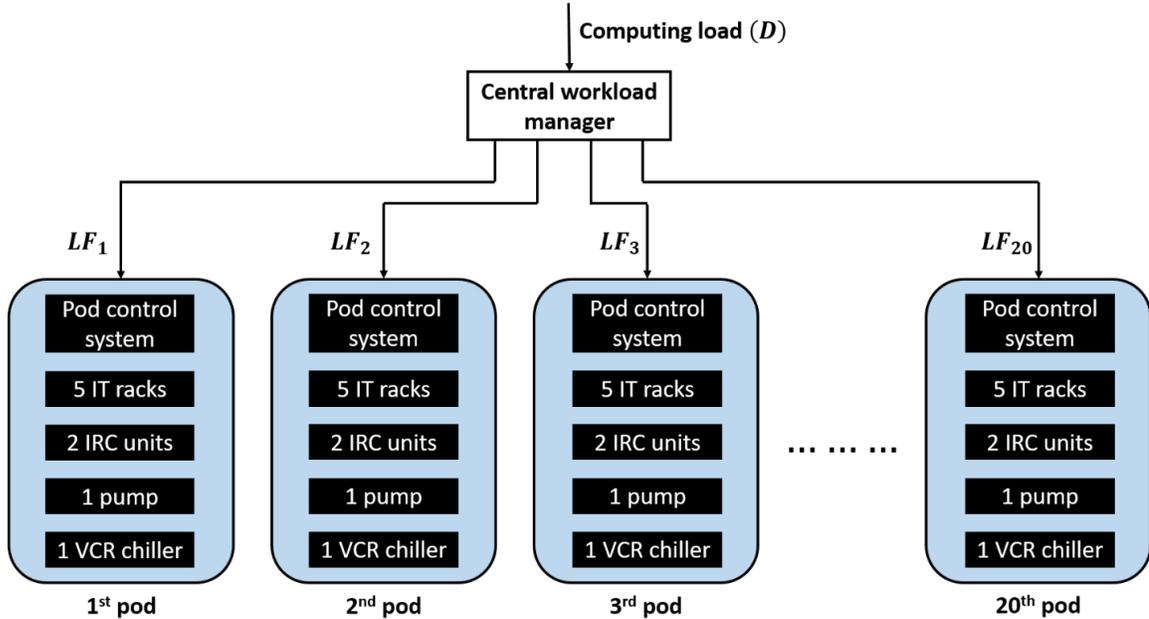


Figure 6.8: Illustration of the pod management strategy for the 100-rack DC considered. The components inside all the pods are identical and are detailed in Figure 6.1a. All pods are connected to a central load management system that presets the LF for each pod.

Consequently, a 100-rack DC infrastructure requires 20 such pods, i.e., $N_p = \frac{N_{total}}{N_r} = 20$. The maximum achievable computing load for this DC $D^{max} = N_p \times N_r \times N_z \times N_s \times u^{max} = 20 \times 5 \times 5 \times 3 \times 1 = 1500$. We consider scenarios with

homogeneous and heterogeneous LF s and distribute the IT load $0 < D \leq 1500$ across the DC pods to determine the influence on CPR_m .

Table 6.11: Results from the heterogeneous vs. homogeneous LF assignment across DC pods to improve CPR_m .

D	Strategy 1 (Heterogeneous LF)	CPR_m^{ht}	Strategy 2 (Homogenous LF)	CPR_m^{hm}	Relative improvement $\left(\frac{CPR_m^{hm} - CPR_m^{ht}}{CPR_m^{ht}}\right) \times 100$
75	$LF=1$ (1 pod) $LF=0$ (19 pods)	0.621	$LF=0.05$ (20 pods)	0.693	11.6%
150	$LF=1$ (2 pods) $LF=0$ (18 pods)	0.621	$LF=0.1$ (20 pods)	0.689	11.0%
300	$LF=1$ (4 pods) $LF=0$ (16 pods)	0.621	$LF=0.2$ (20 pods)	0.682	9.7%
450	$LF=1$ (6 pods) $LF=0$ (14 pods)	0.621	$LF=0.3$ (20 pods)	0.674	8.5%
600	$LF=1$ (8 pods) $LF=0$ (12 pods)	0.621	$LF=0.4$ (20 pods)	0.667	7.3%
750	$LF=1$ (10 pods) $LF=0$ (10 pods)	0.621	$LF=0.5$ (20 pods)	0.659	6.1%
900	$LF=1$ (12 pods) $LF=0$ (8 pods)	0.621	$LF=0.6$ (20 pods)	0.652	4.9%
1050	$LF=1$ (14 pods) $LF=0$ (6 pods)	0.621	$LF=0.7$ (20 pods)	0.644	3.6%
1200	$LF=1$ (16 pods) $LF=0$ (4 pods)	0.621	$LF=0.8$ (20 pods)	0.637	2.4%
1350	$LF=1$ (18 pods) $LF=0$ (2 pods)	0.621	$LF=0.9$ (20 pods)	0.629	1.2%
1500	$LF=1$ (20 pods) $LF=0$ (0 pods)	0.621	$LF=1$ (20 pods)	0.622	0%

The results presented in Table 6.11 reveal that homogeneous LF assignment by the central load manager provides up to 11% relative improvement in CPR_m . The heterogeneous LF management scheme (Strategy 1) only considers $LF = 1$ or 0 for the

Pods with $LF = 0$ do not contribute to computing demand and therefore CPR_m is not a valid metric for them. In contrast, pods with $LF = 1$ utilize the entire computing capacity of the pod, resulting in poor CPR_m . As the value of $D \rightarrow D^{max}$, the relative enhancement in CPR_m using homogeneous LF management diminishes to 0%.

6.5 Conclusion

We have developed an advanced methodology to fulfill multiple objectives determining two-dimensional workload distributions and cooling cycle parameters in a five-rack DC equipped with two IRC units. The methodology combines a high-fidelity, spatiotemporal zonal temperature prediction model while quantifying (1) energy, (2) exergy, and (3) computing efficiencies. Several single- and multi-objective problems are solved, addressing three different objectives leading to joint workload and cooling management in a multi-rack DC. The approach reveals non-trivial trade-offs between cooling cycle efficiencies (energy and exergy) and computing efficiency for a wide range of LF . In contrast to prior investigations that only consider the cooling cycle efficiency for single-rack DCs [1, 21], the present work offers additional benefits by quantifying the degradation (or augmentation) of computing efficiency (measured as in GFLOPs per Watt) for a multi-rack DC. Salient findings include:

- The cooling energy minimization problem populates workload close to the IRC units in a heterogeneous manner and uses a high chilled water temperature $T_{c,w} = 20$ °C.
- Increasing the LF from 0.2 to 0.8 improves the energy-based metric COP_c up to 12.5% and offers significant energy savings.

- Trade-offs are observed across COP_c vs. η_{ex} and COP_c vs. CPR_m , where the desired values for all three objectives cannot be obtained by similar values of cooling parameters and workload distribution.
- The η_{ex} and CPR_m maximization strategy distributes the workload across all of the racks, whereas, the COP_c mediation problem assigns workload close to the IRC units.
- The metric CPR_m is a weak function of the cooling cycle parameters $T_{c,w}$ and \dot{Q}_a . It is predominantly regulated by altering the LF of the DC. Increasing the LF from 0.2 to 0.8 degrades the CPR_m by up to 7.5% resulting in a lower value of GFLOPs at a higher energy expenditure.
- The three-dimensional Pareto fronts and corresponding decision variables help to determine detrimental effects on different objectives.
- An extended approach to manage LF in a large-scale DC is presented to improve the overall CPR_m .

The thermally aware, multi-objective workload management scheme for the five-rack, air-cooled DC reveals essential physical insights and opens avenues for future research in joint thermal and workload management for multi-rack DCs. The above methodology is suitable for offline applications in the DC industry, such as investigating the influence of equipment changes and testing what-if scenarios. For implementation in real-time DC workload and thermal management, a data-driven methodology using a fully connected feedforward multi-layer perceptron (MLP) neural network can be utilized to capture the interplay between Pareto optimal solutions and corresponding decision variables. Given the inputs

(1) $u_{i,j}^z$, (2) $T_{c,w}$, (3) \dot{Q}_a , and (4) LF , the MLP would be able to predict optimal solutions for (1) COP_c , (2) η_{ex} , and (3) CPR_m within a few seconds.

6.6 Appendix: Pressure-flow calculation strategy using the flow network model

The airflow inside a DC enclosure with IRC units is pressure-driven and can be modeled using a flow resistance representation (see Figure 6.2). Our previous experiments demonstrate that the pressure-flow relationship is linear [22]. Hence, the flow resistance is expressed as $R = \Delta P / \dot{Q}$, where ΔP and \dot{Q} are the characteristic pressure drop and flowrate, respectively. The experimentally determined values for different resistances R_H^f , R_V^f , R_H^b , R_V^b , and R_{br} are available in [22]. Different flow interactions of a zone with its neighboring zones are, (1) right $\dot{Q}_{(i+1,j) \rightarrow (i,j)}$, (2) left $\dot{Q}_{(i-1,j) \rightarrow (i,j)}$, (3) top $\dot{Q}_{(i,j+1) \rightarrow (i,j)}$, (4) bottom $\dot{Q}_{(i,j-1) \rightarrow (i,j)}$, (5) leakage flow between front and back chambers through brushes, i.e., $\dot{Q}_{i,j}^{b \rightarrow f}$ or $\dot{Q}_{i,j}^{f \rightarrow b}$, and (6) flow induced by server fans $\dot{Q}_{i,j}^z$. A schematic of these neighboring flow resistances and pressures is provided in Figure 6.9 for a single server zone.

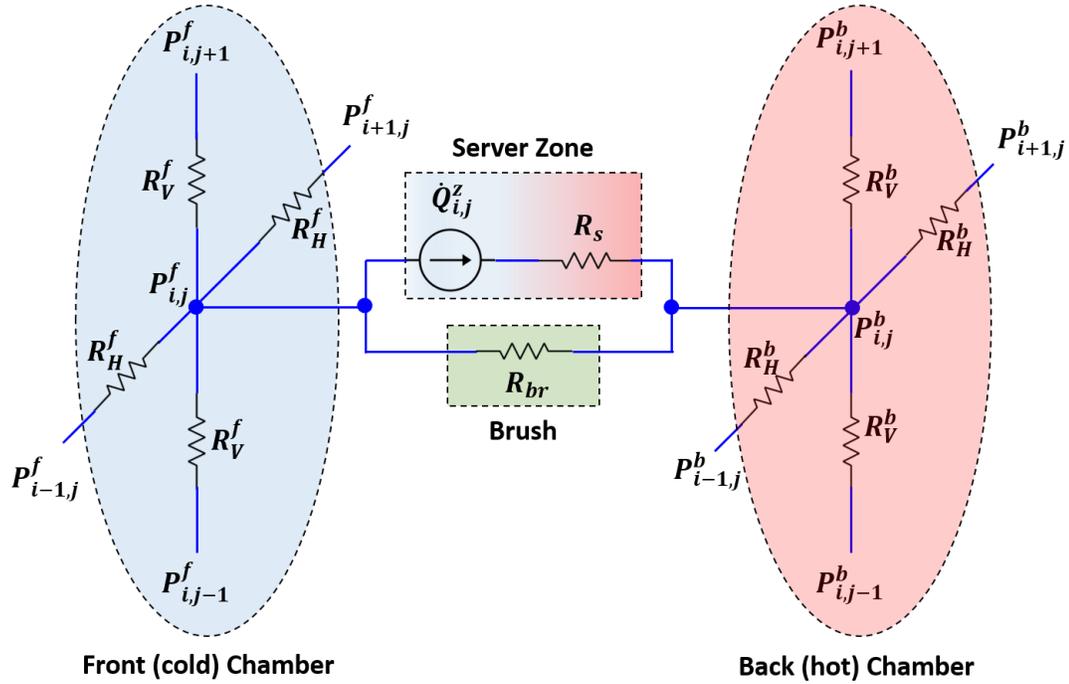


Figure 6.9: Schematic of neighboring flow resistances and pressures for a single server zone. The flow across a resistance depends on the magnitude and sign of the pressure difference. A combination of several such server zones and the IRC units results in the generalized flow network representation shown in Figure 6.2.

Subsequently, the zonal flowrate balance in the front chamber for an active or idle server,

$$\dot{Q}_{(i+1,j) \rightarrow (i,j)}^f + \dot{Q}_{(i-1,j) \rightarrow (i,j)}^f + \dot{Q}_{(i,j+1) \rightarrow (i,j)}^f + \dot{Q}_{(i,j-1) \rightarrow (i,j)}^f + \dot{Q}_{i,j}^{b \rightarrow f} - \dot{Q}_{i,j}^z = 0, \quad (6.1A)$$

$$\Rightarrow \frac{P_{i+1,j}^f - P_{i,j}^f}{R_H^f} + \frac{P_{i-1,j}^f - P_{i,j}^f}{R_H^f} + \frac{P_{i,j+1}^f - P_{i,j}^f}{R_V^f} + \frac{P_{i,j-1}^f - P_{i,j}^f}{R_V^f} + \frac{P_{i,j}^b - P_{i,j}^f}{R_{br}} - \dot{Q}_{i,j}^z = 0, \quad (6.2A)$$

$$\Rightarrow \alpha^f [P_{i+1,j}^f + P_{i-1,j}^f] + \beta^f [P_{i,j+1}^f + P_{i,j-1}^f] + \gamma P_{i,j}^b - (2\alpha^f + 2\beta^f + \gamma) P_{i,j}^f - \dot{Q}_{i,j}^z = 0. \quad (6.3A)$$

While that in the back chamber is expressed as,

$$\dot{Q}_{(i+1,j) \rightarrow (i,j)}^b + \dot{Q}_{(i-1,j) \rightarrow (i,j)}^b + \dot{Q}_{(i,j+1) \rightarrow (i,j)}^b + \dot{Q}_{(i,j-1) \rightarrow (i,j)}^b + \dot{Q}_{i,j}^{f \rightarrow b} + \dot{Q}_{i,j}^z = 0, \quad (6.4A)$$

$$\Rightarrow \frac{P_{i+1,j}^b - P_{i,j}^b}{R_H^b} + \frac{P_{i-1,j}^b - P_{i,j}^b}{R_H^b} + \frac{P_{i,j+1}^b - P_{i,j}^b}{R_V^b} + \frac{P_{i,j-1}^b - P_{i,j}^b}{R_V^b} + \frac{P_{i,j}^f - P_{i,j}^b}{R_{br}} + \dot{Q}_{i,j}^z = 0, \quad (6.5A)$$

$$\Rightarrow \alpha^b [P_{i+1,j}^b + P_{i-1,j}^b] + \beta^b [P_{i,j+1}^b + P_{i,j-1}^b] + \gamma P_{i,j}^f - (2\alpha^b + 2\beta^b + \gamma) P_{i,j}^b + \dot{Q}_{i,j}^z = 0. \quad (6.6A)$$

Likewise, the flowrate balance for the zones in the front chamber adjacent to the IRC units yields,

$$\dot{Q}_{(i+1,j) \rightarrow (i,j)}^f - \frac{\dot{Q}_a}{N_z N_{cu}} = 0, \quad (6.7A)$$

$$\Rightarrow \frac{P_{i+1,j}^f - P_{i,j}^f}{R_H^f} - \frac{\dot{Q}_a}{N_z N_{cu}} = 0, \quad (6.8A)$$

$$\Rightarrow \alpha^f [P_{i+1,j}^f - P_{i,j}^f] - \frac{\dot{Q}_a}{N_z N_{cu}} = 0, \quad (6.9A)$$

and that in the back chamber is expressed as,

$$\dot{Q}_{(i+1,j) \rightarrow (i,j)}^b + \frac{\dot{Q}_a}{N_z N_{cu}} = 0, \quad (6.10A)$$

$$\Rightarrow \frac{P_{i+1,j}^b - P_{i,j}^b}{R_H^b} + \frac{\dot{Q}_a}{N_z N_{cu}} = 0, \quad (6.11A)$$

$$\Rightarrow \alpha^b [P_{i+1,j}^b - P_{i,j}^b] + \frac{\dot{Q}_a}{N_z N_{cu}} = 0. \quad (6.12A)$$

Equations (6.3A), (6.6A), (6.9A), and (6.12A) are the final forms of the pressure-flow calculation, representing two sets of linear equations. The first results in a two-dimensional pressure matrix in the front chamber $[P_{i,j}^f]$, while the latter provides a pressure matrix in the back chamber $[P_{i,j}^b]$. The pressure variation across the geometry is influenced by the following factors.

- The distance between the IRC units and different racks offers a resistance, decreasing the cold air delivered to the racks at farther distances.
- A mismatch between the total air suction of servers and supply by the IRC units creates a pressure difference between the front and back chambers (~ 10 Pa) [35, 38]. According to the sign of the pressure difference, this produces a net leakage flow through the brushes across the two chambers.

- The combined effect of these two influences creates a larger pressure difference between front and back chamber zones which lie farthest from IRC units (R3 in this case).

6.7 References

- [1] S. MirhoseiniNejad, H. Moazamigoodarzi, G. Badawy, and D. G. Down, "Joint data center cooling and workload management: A thermal-aware approach," *Future Generation Computer Systems*, vol. 104, pp. 174-186, 2020.
- [2] S. MirhoseiniNejad, G. Badawy, and D. G. Down, "Holistic thermal-aware workload management and infrastructure control for heterogeneous data centers using machine learning," *Future Generation Computer Systems*, vol. 118, pp. 208-218, 2021.
- [3] B. Shi and A. Srivastava, "Unified datacenter power management considering on-chip and air temperature constraints," *Sustainable Computing: Informatics and Systems*, vol. 1, pp. 91-98, 2011.
- [4] S. K. Garg, S. K. Gopalaiyengar, and R. Buyya, "SLA-based resource provisioning for heterogeneous workloads in a virtualized cloud datacenter," in *International conference on Algorithms and architectures for parallel processing*, 2011, pp. 371-384.
- [5] S. Ghorbani, Z. Yang, P. B. Godfrey, Y. Ganjali, and A. Firoozshahian, "Drill: Micro load balancing for low-latency data center networks," in *Proceedings of the Conference of the ACM Special Interest Group on Data Communication*, 2017, pp. 225-238.

- [6] R. Gupta, S. Asgari, H. Moazamigoodarzi, S. Pal, and I. K. Puri, "Cooling architecture selection for air-cooled Data Centers by minimizing exergy destruction," *Energy*, vol. 201, p. 117625, 2020.
- [7] H. Moazamigoodarzi, P. J. Tsai, S. Pal, S. Ghosh, and I. K. Puri, "Influence of cooling architecture on data center power consumption," *Energy*, vol. 183, pp. 525-535, 2019.
- [8] K. Ebrahimi, G. F. Jones, and 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*, vol. 31, pp. 622-638, 2014.
- [9] Z. Ding, Y. Cao, L. Xie, Y. Lu, and P. Wang, "Integrated stochastic energy management for data center microgrid considering waste heat recovery," *IEEE Transactions on Industry Applications*, vol. 55, pp. 2198-2207, 2019.
- [10] P. Huang, B. Copertaro, X. Zhang, J. Shen, I. Löfgren, M. Rönnelid, *et al.*, "A review of data centers as prosumers in district energy systems: Renewable energy integration and waste heat reuse for district heating," *Applied Energy*, vol. 258, p. 114109, 2020.
- [11] Y. Bai, L. Gu, and X. Qi, "Comparative study of energy performance between chip and inlet temperature-aware workload allocation in air-cooled data center," *Energies*, vol. 11, p. 669, 2018.
- [12] Q. Tang, S. K. S. Gupta, and G. Varsamopoulos, "Energy-efficient thermal-aware task scheduling for homogeneous high-performance computing data centers: A

cyber-physical approach," *IEEE Transactions on Parallel and Distributed Systems*, vol. 19, pp. 1458-1472, 2008.

- [13] J. D. Moore, J. S. Chase, P. Ranganathan, and R. K. Sharma, "Making Scheduling" Cool": Temperature-Aware Workload Placement in Data Centers," in *USENIX annual technical conference, General Track*, 2005, pp. 61-75.
- [14] R. K. Sharma, C. E. Bash, C. D. Patel, R. J. Friedrich, and J. S. Chase, "Balance of power: Dynamic thermal management for internet data centers," *IEEE Internet Computing*, vol. 9, pp. 42-49, 2005.
- [15] M. T. Chaudhry, T. Ling, S. A. Hussain, and A. Manzoor, "Minimizing thermal stress for data center servers through thermal-aware relocation," *The Scientific World Journal*, vol. 2014, 2014.
- [16] L. Silva-Llanca, A. Ortega, K. Fouladi, M. del Valle, and V. Sundaralingam, "Determining wasted energy in the airside of a perimeter-cooled data center via direct computation of the Exergy Destruction," *Applied Energy*, vol. 213, pp. 235-246, 2018.
- [17] A. Bhalerao, K. Fouladi, L. Silva-Llanca, and A. P. Wemhoff, "Rapid prediction of exergy destruction in data centers due to airflow mixing," *Numerical Heat Transfer, Part A: Applications*, vol. 70, pp. 48-63, 2016.
- [18] A. J. Shah, V. P. Carey, C. E. Bash, and C. D. Patel, "Exergy analysis of data center thermal management systems," *Journal of Heat Transfer*, vol. 130, 2008.

- [19] A. J. Díaz, R. Cáceres, J. M. Cardemil, and L. Silva-Llanca, "Energy and exergy assessment in a perimeter cooled data center: The value of second law efficiency," *Applied Thermal Engineering*, vol. 124, pp. 820-830, 2017.
- [20] K. Fouladi, A. P. Wemhoff, L. Silva-Llanca, K. Abbasi, and A. Ortega, "Optimization of data center cooling efficiency using reduced order flow modeling within a flow network modeling approach," *Applied Thermal Engineering*, vol. 124, pp. 929-939, 2017.
- [21] R. Gupta, H. Moazamigoodarzi, S. MirhoseiniNejad, D. G. Down, and I. K. Puri, "Workload management for air-cooled data centers: An energy and exergy based approach," *Energy*, vol. 209, p. 118485, 2020.
- [22] H. Moazamigoodarzi, R. Gupta, S. Pal, P. J. Tsai, S. Ghosh, and I. K. Puri, "Modeling temperature distribution and power consumption in IT server enclosures with row-based cooling architectures," *Applied Energy*, vol. 261, p. 114355, 2020.
- [23] S. J. Ovaska, R. E. Dragseth, and S. A. Hanssen, "Direct-to-chip liquid cooling for reducing power consumption in a subarctic supercomputer centre," *International Journal of High Performance Computing and Networking*, vol. 9, pp. 242-249, 2016.
- [24] A. A. Moskovsky, E. A. Druzhinin, A. B. Shmelev, V. V. Mironov, and A. Semin, "Server Level Liquid Cooling: Do Higher System Temperatures Improve Energy Efficiency?," *Supercomputing frontiers and innovations*, vol. 3, pp. 67-74, 2016.

- [25] D. C. Price, M. A. Clark, B. R. Barsdell, R. Babich, and L. J. Greenhill, "Optimizing performance-per-watt on GPUs in high performance computing," *Computer Science-Research and Development*, vol. 31, pp. 185-193, 2016.
- [26] J. Shin and E. Seo, "Impact of Data Center Cooling Technology to Effectiveness of Turbo-Mode," in *2020 International Conference on Information and Communication Technology Convergence (ICTC)*, 2020, pp. 1691-1693.
- [27] Y. Li, X. Wang, P. Luo, and Q. Pan, "Thermal-aware hybrid workload management in a green datacenter towards renewable energy utilization," *Energies*, vol. 12, p. 1494, 2019.
- [28] Z. Abbasi, G. Varsamopoulos, and S. K. Gupta, "Tacoma: Server and workload management in internet data centers considering cooling-computing power trade-off and energy proportionality," *ACM Transactions on Architecture and Code Optimization (TACO)*, vol. 9, pp. 1-37, 2012.
- [29] X. Li, P. Garraghan, X. Jiang, Z. Wu, and J. Xu, "Holistic virtual machine scheduling in cloud datacenters towards minimizing total energy," *IEEE Transactions on Parallel and Distributed Systems*, vol. 29, pp. 1317-1331, 2017.
- [30] Q. Fang, J. Wang, Q. Gong, and M. Song, "Thermal-aware energy management of an HPC data center via two-time-scale control," *IEEE Transactions on Industrial Informatics*, vol. 13, pp. 2260-2269, 2017.
- [31] Z. Liu, Y. Chen, C. Bash, A. Wierman, D. Gmach, Z. Wang, *et al.*, "Renewable and cooling aware workload management for sustainable data centers," in *Proceedings of the 12th ACM SIGMETRICS/PERFORMANCE joint international*

conference on Measurement and Modeling of Computer Systems, 2012, pp. 175-186.

- [32] S. Mirhoseininejad, G. Badawy, and D. G. Down, "A Data-driven, Multi-setpoint Model Predictive Thermal Control System for Data Centers," *Journal of Network and Systems Management*, vol. 29, pp. 1-22, 2021.
- [33] S. Asgari, S. MirhoseiniNejad, H. Moazamigoodarzi, R. Gupta, R. Zheng, and I. K. Puri, "A gray-box model for real-time transient temperature predictions in data centers," *Applied Thermal Engineering*, vol. 185, p. 116319, 2021.
- [34] S. Asgari, H. Moazamigoodarzi, P. J. Tsai, S. Pal, R. Zheng, G. Badawy, *et al.*, "Hybrid surrogate model for online temperature and pressure predictions in data centers," *Future Generation Computer Systems*, vol. 114, pp. 531-547, 2021.
- [35] H. Moazamigoodarzi, S. Pal, S. Ghosh, and I. K. Puri, "Real-time temperature predictions in it server enclosures," *International Journal of Heat and Mass Transfer*, vol. 127, pp. 890-900, 2018.
- [36] RITTAL, "C.W. TopTherm LCP rack/inline, Model number: 3311.130," 2015.
- [37] TRANE, "Air-cooled liquid chillers, 10 to 60 Tons," ed, 2004.
- [38] H. Moazamigoodarzi, S. Pal, D. Down, M. Esmalifalak, and I. K. Puri, "Performance of a rack mountable cooling unit in an IT server enclosure," *Thermal Science and Engineering Progress*, vol. 17, p. 100395, 2020.
- [39] J. W. VanGilder, Z. Pardey, C. Healey, and X. Zhang, "A compact server model for transient data center simulations," *ASHRAE Transactions*, vol. 119, pp. 358-370, 2013.

- [40] C. Jin, X. Bai, C. Yang, W. Mao, and X. Xu, "A review of power consumption models of servers in data centers," *Applied Energy*, vol. 265, p. 114806, 2020.
- [41] S.-W. Ham, M.-H. Kim, B.-N. Choi, and J.-W. Jeong, "Simplified server model to simulate data center cooling energy consumption," *Energy and Buildings*, vol. 86, pp. 328-339, 2015.
- [42] H. Packard, "Quickspecs: HP ProLiant DL360 Generation 5 (G5)," in *Product datasheet*, 27, ed, 2007, pp. 1-43.
- [43] J. M. Gordon, K. Ng, and H. Chua, "Optimizing chiller operation based on finite-time thermodynamics: universal modeling and experimental confirmation," *International Journal of Refrigeration*, vol. 20, pp. 191-200, 1997.
- [44] S. M. M. Nejad, G. Badawy, and D. G. Down, "EAWA: Energy-aware workload assignment in data centers," in *2018 International Conference on High Performance Computing & Simulation (HPCS)*, 2018, pp. 260-267.
- [45] R. Steinbrecher and R. Schmidt, "Data center environments," *ASHRAE J*, vol. 42, p. e49, 2011.

7 Internal waste heat utilization from heterogeneous data centers using adsorption chillers

This chapter is reproduced from “*Waste heat recovery in a data center with an adsorption chiller: Technical and economic analysis*”, **Rohit Gupta** and Ishwar K. Puri, *Energy Conversion and Management (2021): 114576 (In press)*, Copyright © 2021, with permission from Elsevier. The author of this thesis is the first author and the primary contributor to this publication.

7.1 Abstract

Data center and high-performance computing cluster computing load have increased significantly over the past decade, resulting in higher energy consumption and a larger carbon footprint. We provide a method to utilize the waste heat from a data center room that contains both air- and water-cooled racks. Heat dissipated by the water-cooled racks drives a silica gel-H₂O adsorption chiller, which produces chilled water required for in-row cooling units installed in the air-cooled racks. Hence, the cooling load of the existing vapor-compression chiller is either partially or fully eliminated, offering significant energy savings. Thermodynamic, economic, and environmental benefits of the proposed strategy are compared against a baseline infrastructure without waste heat utilization. Parametric sensitivity analyses reveal that the load sharing factor, inlet temperature of water-cooled racks, chilled water setpoint, airflow of the in-row cooling units, and the re-cooled water temperature all influence the energy savings efficiency, annual variable costs, CO₂ emission reduction, and cost payback period. In this manner, energy savings efficiency is

enhanced by 22.5% and annual CO₂ emissions reduced by 104 tons. The cost payback periods with 30 and 90 kW adsorption chiller retrofits are as low as 379 and 285 days. Raising the water-cooled inlet temperature enhances the energy savings efficiency by as much as 16% but that also degrades computing performance per watt by 6%.

7.2 Introduction

With increasingly intensive and ubiquitous computing, the data center (DC) and high-performance cluster (HPC) have been significantly amplified, which account for $\approx 4\%$ of worldwide electricity usage [1]. The COVID-19 pandemic has increased virtualization, raising the average load on DCs by 30% [2]. Based on the architecture employed, cooling consumes 30-40% of DC energy, which leads to significant energy waste and higher operational expenditure (OPEX) [3, 4].

Air-cooling is a popular and convenient choice for DC and HPC applications. Modular air-cooled systems, such as those employing in-row, rear-door, and rack-mountable cooling, typically manage a maximum heat load of 20 kW heat per rack [5-9]. Since the heat transfer capacity of air is limited, improvements in heat transfer have been sought through liquid cooling. Newer liquid-cooled servers use water-cooled cold plates, micro-channel heat sinks, or enclosed spray-cooled heat sinks to cool CPUs and GPUs. Examples include commercially available servers such as IBM Power System AC922, Lenovo ThinkSystem SD650, DELL PowerEdge C6525 with CoolIT liquid-cooled heat sink, and HPE SGI 8600 [10-12]. However, liquid cooling for other heat-dissipating components, e.g., for dual-inline memory module (DIMM), solid-state drives (SSD), and

PCIe cards, in a server chassis is yet unavailable [13, 14]. Immersion cooling immerses the entire server chassis in an electrically non-conducting coolant, such as a dielectric fluid but has not been broadly adopted due to skepticism about its maintenance [15]. Hence, water-cooled cold plates are preferred in liquid-cooled servers.

The transition to liquid cooling poses two challenges. First, replacing an entire air-cooled legacy server room with liquid-cooled racks requires a significantly high capital expenditure (CAPEX), necessitating an extended payback period [16, 17]. Second, since commercially available liquid-cooled servers utilize liquid-cooling only for CPUs and GPUs, additional air-cooling must take care of the remaining 30-35% of total heat dissipation [13, 14]. We consider this infrastructure heterogeneity and develop a solution to minimize overall OPEX for such systems. High-density liquid-cooled racks dissipate up to 50-60 kW of heat per rack through the water stream. In a closed-loop system, this heat is typically rejected by cooling towers and condensers. Waste heat can be extracted from the exhaust hot water of liquid-cooled servers, but this is challenging at relatively low temperatures of 40-70 °C [10].

Prior investigations of DC waste heat extraction include (a) hot water production for domestic and industrial applications [10, 18], (b) space heating in district energy networks [19-25], (c) an organic Rankine cycle (ORC) driven on-site electricity production [26-28], (d) preheating boiler feedwater for a DC-located power plant [29], and (e) producing chilled water with a LiBr-H₂O absorption chiller [30-32]. However, these applications have drawbacks.

First, community water or space heating networks are typically located some distance away from the DC. Therefore, transporting low-quality waste heat (at 40-70 °C) to a district heating (DH) network without boosting the temperature, e.g., using heat pumps, is infeasible. Besides, the inclusion of the heat pumps increases the CAPEX and OPEX of the heat recovery system [19]. Second, the safe operating temperatures for computing chips lie below 85 °C [33], where the first-law efficiency of an ORC is approximately 5% with a prolonged payback period of 5-8 years [26-28]. This makes ORC-based waste heat recovery unattractive for DC administrators unless they are offered electricity price incentives. Finally, a LiBr-H₂O or NH₃-H₂O based absorption refrigeration system with significantly high waste heat temperatures in the 65-100 °C range has been proposed [30-32]. Below 65°C, this system shuts down and has a poor coefficient of performance (COP), making it infeasible for integration with a DC. Absorption chillers also have a low power to weight ratio, which is incompatible considering the constrained space in enterprise DCs. Moreover, the absorbents used in these systems are highly corrosive and environmentally unsustainable.

Instead, we propose a strategy where an adsorption chiller, with silica gel as the adsorbent material extracts waste heat from the hot water streams produced by liquid-cooled racks. The hot water drives the adsorption refrigeration process and produces chilled water. Our conceptual DC consists of both liquid-cooled and air-cooled racks. The air-cooled racks have in-row cooling units (IRCUs) that use chilled water to extract heat. This is simulated using a benchmarked reduced-order model from our previous work [6]. Two sources of chilled water are considered for the air-cooled racks, one produced by an

adsorption chiller and the other with a vapor-compression refrigeration (VCR) chiller. The aim is to minimize overall OPEX by utilizing the adsorption chiller most of the time. Unlike ORC and DH networks, our approach utilizes waste heat internally within a DC, requiring minimal retrofit that results in a shorter payback period.

With this waste heat utilization strategy, we (1) assess the energy, economic and environmental impact of our solution; (2) examine the sensitivity of key parameters that augment (or lower) system performance; (3) demonstrate the applicability of the solution for DCs with different heat-generation capacities; and (4) elucidate the tradeoffs between computing performance per watt of liquid-cooled racks and waste heat recovery efficiency.

7.3 Description of the data center cooling infrastructure with waste heat recovery

The technical and economic efficiency of a hybrid DC infrastructure that contains water-cooled high-density and air-cooled low-density computing racks can be enhanced by the internal utilization of DC waste heat. The original system (OS) consists of water-cooled HPC racks and air-cooled server racks (see Figure 7.1a). Details of the heat load and cooling infrastructure of these systems are based on a DC located in Toronto, Canada. The maximum heat loads of the water-cooled and air-cooled racks are 350 kW and 50 kW, respectively. The water-cooled server racks are cooled with intermediate rack-based water-water heat exchangers in the OS that has a chiller-less building water delivery system. The air-cooled racks utilize two IRCUs that receive cold water from a VCR chiller that has an ambient air-cooled condenser (see Figure 7.1a).

The new system (NS) that we design is shown in Figure 7.1b. It connects the two types of racks using a waste heat recovery system consisting of (1) an adsorption chiller, (2) a water-water waste heat recuperator, (3) a wet cooling tower, (4) cooling tower fans, and (5) circulating pumps. The inlet water stream to the liquid-cooled racks becomes warmer at the server exhaust due to heat transfer. Aggregated warm water streams from all server exhausts then reject a fraction of their heat, which would have been otherwise wasted, to the adsorption chiller through an intermediate heat exchanger. This drives adsorption in the chiller and produces the chilled water required by the IRCUs located in air-cooled racks. Thus, part of the cooling load previously handled by the VCR chiller is now moved to the adsorption chiller, reducing VCR chiller energy consumption. The performance of the adsorption chiller is degraded by lowering the temperature of the waste heat obtained from the liquid-cooled server exhausts. Since heat dissipation from both liquid- and air-cooled racks varies based on end-user computing demand, having both the adsorption and VCR chillers in the infrastructure is necessary.

The following assumptions were incorporated to simplify the thermodynamic simulation of the data center: (1) Transient fluctuation of temperature, flowrates, and heat loads are neglected, (2) The effectiveness of the waste heat exchanger was assumed constant, (3) The part-load performance of the adsorption and vapor compression chillers was captured using a regression-based model, which neglected the internal refrigeration mechanisms, (4) Each of the equipment was treated as control volumes with energy balance, (5) Instead of detailed modelling of building chilled water system associated with liquid-cooled infrastructure, its performance is based on a constant coefficient of

performance (COP) obtained from the literature. Intricate system-level modelling of each liquid-cooled rack was not adopted. Instead, the liquid-cooled racks were considered as heat sources producing hot water at a certain temperature.

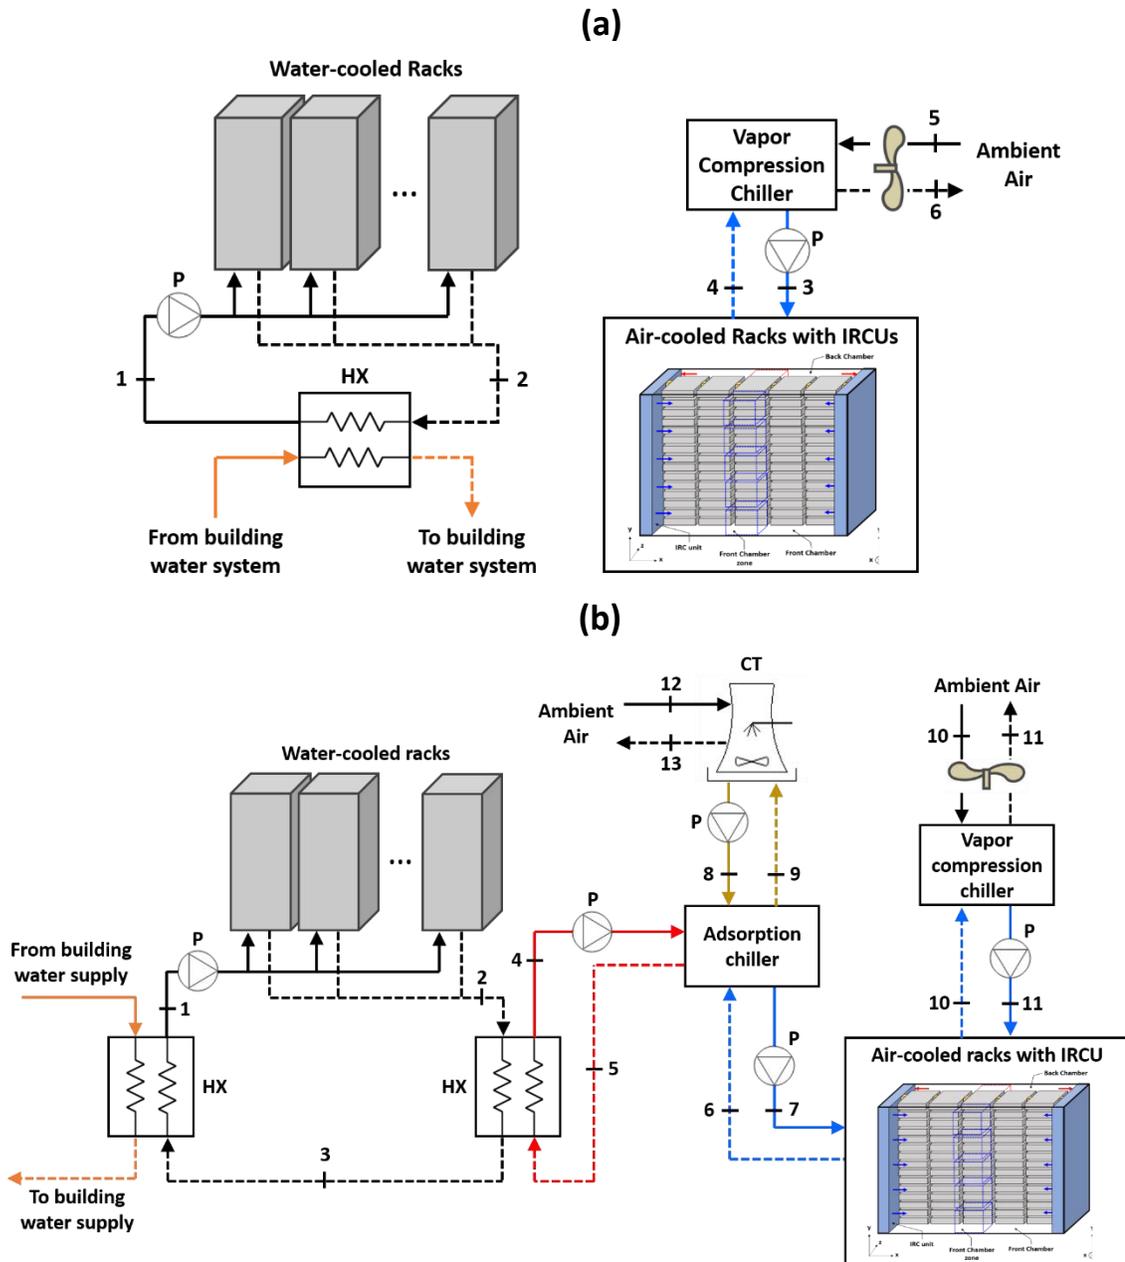


Figure 7.1: System layout for a DC with both liquid-cooled and air-cooled racks. (a) The original system (OS), where the liquid-cooled and air-cooled systems are decoupled, and

(b) the new system (NS), where these two systems are coupled through an adsorption chiller.

7.4 Mathematical models

7.4.1 Liquid-cooled infrastructure

The liquid-cooled racks contain servers that use water-cooled cold plates mounted on the central processing units (CPUs). Within the DC, there are 10 liquid-cooled racks, each having 40 servers and a rack-based water-water heat exchanger. The liquid-cooled cold plates reject heat to the inlet water supplied by a building cooling system. In the OS, the heat transported by the warm water is rejected to the environment through a chiller-less building cooling system (points 1 and 2 in Figure 7.1a). In the NS, this heat is partially recovered through an intermediate heat exchanger or recuperator (points 2 and 3 in Figure 7.1b). We assume that the entire electrical energy input to a liquid-cooled server is rejected to water. Consequently, the following steady-state energy balance provides the exhaust warm water temperature,

$$T_{w,o}^{lc} = T_{w,i}^{lc} + \frac{Q_{lc}}{\rho_w C_{p,w} \dot{V}_{lc}}, \quad (7.1)$$

where $T_{w,i}^{lc}$ and $T_{w,o}^{lc}$ water inlet and outlet temperatures, $Q_{lc} = \sum_k \sum_j Q_{j,k}^{lc}$ the sum of heat dissipation of all the liquid-cooled servers, $\dot{V}_{lc} = \sum_k \sum_j \dot{V}_{j,k}^{lc}$ the sum of waterflow for all servers [34], ρ_w the density of water, and $C_{p,w}$ the specific heat capacity of water. The average CPU temperature of the liquid-cooled servers as a function of $T_{w,i}^{lc}$ is obtained from the literature as follows [35, 36],

$$T_{CPU}^{lc} = A_1 + A_2 T_{w,i}^{lc}, \quad (7.2)$$

where $A_1 = 28.81$ and $A_2 = 0.88$ are valid for server inlet water temperatures in the $20 \leq T_{w,i}^{lc}(\text{°C}) \leq 70$ range.

To perform the technical and economic analysis, we calculate the power consumption and OPEX for a building cooling system consisting of (a) a building cooling tower, (b) fans, and (c) pumping system. Instead of a detailed analysis, the performance of a liquid-cooled DC located in Toronto is obtained from the literature [25] (see Table 7.1). The total power consumption for the building cooling system is,

$$\dot{P}_{lc} = \frac{Q_{lc}}{\text{COP}_{lc}} = \frac{\sum_k \sum_j Q_{j,k}^{lc}}{\text{COP}_{lc}}. \quad (7.3)$$

Table 7.1: Performance and variable cost data for the building cooling tower used to extract heat generated by the liquid-cooled DC [25, 34].

Quantity	Value
Number of liquid-cooled racks (k)	10
Number of server in each rack (j)	40
Design water flowrate of servers ($\dot{V}_{j,k}^{lc}$)	0.264 GPM ($1.67 \times 10^{-5} \text{ m}^3 \text{ s}^{-1}$)
Maximum total water flowrate of servers (\dot{V}_{lc})	105 GPM ($6.68 \times 10^{-3} \text{ m}^3 \text{ s}^{-1}$)
Coefficient of performance (COP_{lc})	17.5
Water cost (c_{water})	0.008 CA\$/Ton-h (2.27×10^{-3} CA\$/kW ref-h)
Chemical cost (c_{chem})	0.001 CA\$/Ton-h (2.84×10^{-4} CA\$/kW ref-h)

7.4.2 Air-cooled infrastructure

7.4.2.1 Server racks with in-row cooling units

The air-cooled infrastructure with five racks and two IRCUs (see Figure 7.2b) is simulated with a previously experimentally validated flow network model (FNM) [6]. The five-rack system houses 125 1U servers, which generate a maximum of 50 kW of heat. The FNM

adopts a zonal approach to predict temperatures at various locations within air-cooled DC. Essential inputs to the model are (1) total airflow \dot{V}_a^{IRCU} and waterflow \dot{V}_w^{IRCU} of IRCUs, (2) chilled water temperature entering the IRCUs $T_{w,i}^{IRCU}$ and (3) total heat load distributed among all servers $Q_{ac} = \sum_k \sum_j Q_{j,k}^{ac}$. Given these inputs, the model yields (1) cold (front) chamber $T_{j,k}^f$ and hot (back) chamber $T_{j,k}^b$ air temperature distributions, (2) warm water temperature of IRCUs $T_{w,o}^{IRCU}$, (3) mean cold air supply $T_{a,o}^{IRCU}$ and hot air return $T_{a,i}^{IRCU}$ temperatures of the IRCUs, (4) power consumption of the fans inside IRCUs \dot{P}_{IRCU} . Figure 7.2a summarizes the input and output parameters of the FNM for the air-cooled racks [6].

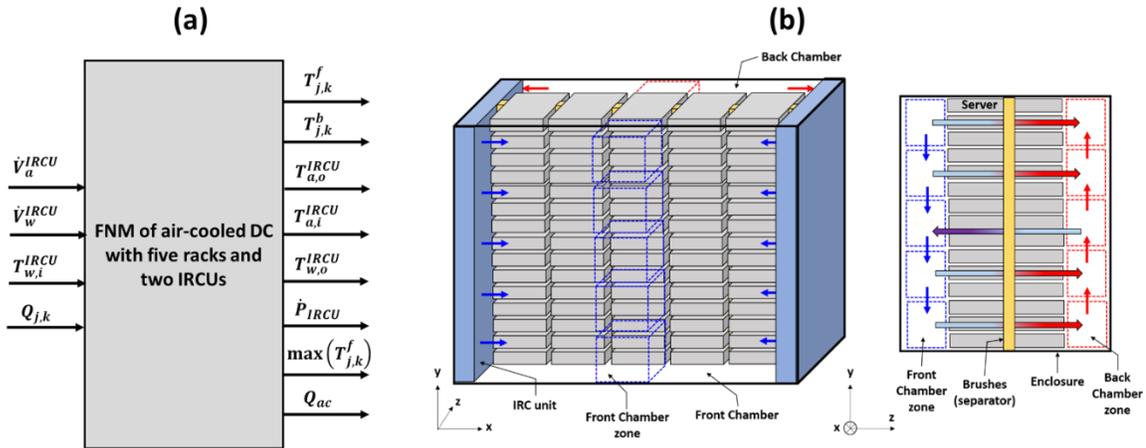


Figure 7.2: (a) Input and output parameters of the FNM used for simulating the air-cooled infrastructure with five server racks and two in-row cooling units (IRCUs). (b) Rack geometry showing the arrangement of servers and IRCUs [6].

The heat Q_{ac} generated by the air-cooled server racks in the NS is dissipated with chilled water at the specified temperature $T_{w,i}^{IRCU}$ obtained from two sources, (1) the existing VCR chiller and (2) the newly added adsorption chiller, leading to the following energy balance,

$$Q_{ac} = (1 - \alpha)Q_e^{ads} + \alpha Q_e^{vcr}, \quad (7.4)$$

where Q_e^{ads} and Q_e^{vcr} are the cooling loads on the evaporators of adsorption and VCR chillers, respectively, and α is the load sharing factor that varies between 0 to 1. Since the adsorption chiller is absent from the OS, $Q_{ac} = Q_e^{vcr}$.

7.4.2.2 Vapor compression chiller

For both OS and NS, the air-cooled infrastructure is integrated with a VCR chiller, which supplies the chilled water required for the IRCUs (points 4 and 3 in Figure 7.1a, and points 10 and 11 in Figure 7.1b). Since electrical power is required to drive the chiller compressor, the power consumption under part-load scenarios is required. Based on a previous application [6], the well-known Gordon-Ng single equation model for VCR chiller is used [37],

$$\left[\left(\frac{1}{\text{COP}_{vcr}} + 1 \right) - \frac{T_{c,i}^{vcr}}{T_{e,o}^{vcr}} \right] Q_e^{vcr} = -B_1 + B_2 T_{c,i}^{vcr} - B_3 (T_{c,i}^{vcr} / T_{e,o}^{vcr}), \quad (7.5)$$

where the coefficient of performance $\text{COP}_{vcr} = Q_e^{vcr} / \dot{P}_{vcr}$, Q_e^{vcr} the heat load on the chiller evaporator, i.e., the fraction of cooling load of air-cooled server racks, $T_{e,o}^{vcr}$ the chilled water temperature generated at the chiller evaporator, $T_{c,i}^{vcr}$ the ambient air temperature entering the VCR chiller condenser, and \dot{P}_{vcr} the power consumption of the chiller in the compressor and water circulating pump. All temperatures are expressed in K, and \dot{P}_{vcr} and Q_e^{vcr} are in kW. To evaluate the model constants B_i , the performance data of a commercially available VCR chiller is used [38]. The operating parameters and resulting

performance constants for VCR chillers with two different cooling capacities are presented in Table 7.2. An energy balance on the outer control volume of the VCR chiller yields,

$$\dot{P}_{vcr} + \rho_w C_{p,w} \dot{V}_e^{vcr} (T_{e,i}^{vcr} - T_{e,o}^{vcr}) + \rho_a C_{p,a} \dot{V}_c^{vcr} (T_{c,i}^{vcr} - T_{c,o}^{vcr}) = 0, \quad (7.6)$$

where ρ_a denotes the density of air, $C_{p,a}$ the specific heat capacity of air, \dot{V}_c^{vcr} the airflow of the VCR chiller condenser fans. $T_{e,i}^{vcr}$ the water temperature entering the chiller evaporator, and $T_{c,o}^{vcr}$ the exhaust air temperature of the chiller condenser.

Table 7.2: Specifications and regression constants for VCR chillers with two different cooling capacities [38].

Model name	Nominal cooling capacity (kW)	Allowable range for \dot{V}_e^{vcr} (GPM)	\dot{V}_c^{vcr} (CFM)	Model constants		
				B_1	B_2	B_3
TRANE CGA-120 (60 Hz)	35.2	12 to 36	8120	26.16	0.2	24.9
TRANE CGAF-C25 (60 Hz)	87.9	30 to 90	21650	64.38	0.5	66.0

7.4.3 Heat recovery system

7.4.3.1 Adsorption chiller

In contrast to a conventional VCR chiller, an adsorption chiller utilizes low-temperature waste heat, which would have otherwise been wasted, to drive the refrigeration process. It consists of three coolant loops for (1) high, (2) medium, and (3) low -temperatures (see Figure 7.1b). The high-temperature (HT) loop connects the adsorption chiller generator to the waste heat source, i.e., the exhaust water stream of the liquid-cooled racks (points 4 and 5 in Figure 7.1b). The medium-temperature (MT) loop couples the adsorption chiller condenser with a wet cooling tower, providing necessary heat rejection (points 8 and 9 in Figure 7.1b). The third loop is on the low temperature (LT) side at the chiller evaporator,

where chilled water is produced (points 6 and 7 in Figure 7.1b). In this case, the chilled water is supplied to the IRCUs installed in the air-cooled racks.

The off-design performance for an adsorption chiller is simulated using the Gordon-Ng single equation model for thermally-driven chillers [37],

$$\left[\frac{1}{\text{COP}_{ads}} - \left(\frac{T_{c,i}^{ads} - T_{e,o}^{ads}}{T_{e,o}^{ads}} \right) \left(\frac{T_{g,i}^{ads}}{T_{g,i}^{ads} - T_{c,i}^{ads}} \right) \right] = \frac{C_1 - C_2 \frac{T_{c,i}^{ads}}{T_{g,i}^{ads}}}{\left[Q_e^{ads} \left(\frac{T_{g,i}^{ads} - T_{e,i}^{ads}}{T_{g,i}^{ads}} \right) \right]} \quad (7.7)$$

where $\text{COP}_{ads} = Q_e^{ads} / Q_g^{ads}$, Q_e^{ads} represents the cooling load on the chiller evaporator, Q_g^{ads} the waste heat added to the generator side, $T_{g,i}^{ads}$ the inlet temperature of hot water to HT side, $T_{e,i}^{ads}$ and $T_{e,o}^{ads}$ the inlet and outlet temperatures at the LT loop, and $T_{c,i}^{ads}$ the inlet water temperature to the MT side supplied by the wet cooling tower. To calculate the performance constants C_1 and C_2 , performance datasheets of the InvenSor adsorption are used [39, 40]. The adsorption chiller can only produce chilled water for generator inlet temperatures $T_{g,i}^{ads} \geq 40$ °C. Essential operating parameters and resulting performance constants for adsorption chillers with two different cooling capacities are shown in Table 7.3. An energy balance on the outer boundary of the adsorption chiller yields,

$$Q_c^{ads} = Q_g^{ads} + Q_e^{ads} = \rho_w \dot{V}_w^{ads} C_{p,w} (T_{c,o}^{ads} - T_{c,i}^{ads}), \quad (7.8)$$

where $T_{c,o}^{ads}$ denotes the water temperature leaving the adsorption chiller condenser, \dot{V}_w^{ads} the waterflow of condenser, and Q_c^{ads} the heat to be rejected by the cooling tower coupled to the adsorption chiller condenser.

Table 7.3: Specifications and regression constants for adsorption chiller with two different cooling capacities [39, 40].

Model Name	Nominal cooling capacity (kW)	Power consumption \dot{P}_{ads} (kW)	\dot{V}_e^{ads} (GPM)	\dot{V}_c^{ads} (GPM)	\dot{V}_g^{ads} (GPM)	Model constants	
						C_1	C_2
InvenSor LTC30 e-plus	29.5	0.895	29.0	50.2	27.7	70.5	75.3
InvenSor LTC90 e-plus	88.5	2.685	87.2	150.6	83.2	221.3	225.8

7.4.3.2 Waste heat recuperator

The waste heat rejected by the liquid-cooled racks at the temperature $T_{w,o}^{lc}$ is provided to the water stream leaving the adsorption chiller generator at $T_{g,o}^{ads}$ through an intermediate heat exchanger (points 2, 3, 4, and 5 in Figure 7.1b). This heat exchange upgrades the generator side temperature from $T_{g,o}^{ads}$ to $T_{g,i}^{ads}$ and drives adsorption cooling. The value of $T_{g,i}^{ads}$ is determined using the effectiveness ε_{hex} for the heat exchanger,

$$\varepsilon_{hex} = \frac{\dot{V}_g^{ads}(T_{g,i}^{ads} - T_{g,o}^{ads})}{\min[\dot{V}_g^{ads}, \dot{V}_{lc}] \times (T_{w,o}^{lc} - T_{g,o}^{ads})}. \quad (7.9)$$

To simplify the thermodynamic analysis, we assume $\varepsilon_{hex} = 0.9$ for the water-water waste heat recuperator [32]. Values for \dot{V}_g^{ads} and $\dot{V}_{lc} = \sum_k \sum_j \dot{V}_{j,k}^{lc}$ are provided in Table 7.3 and Table 7.1, respectively.

7.4.3.3 Cooling tower and supporting systems

To obtain the necessary heat rejection for the adsorption chiller condenser Q_c^{ads} , consistent with the DC literature, a counter-flow mechanical draft wet cooling tower is modelled using the effectiveness-NTU method [41-43]. This approach considers the heat and mass transfers of the cooling tower similar to that of a single-phase heat exchanger. The heat rejected by the cooling tower,

$$Q_{ct} = \rho_a \dot{V}_a^{ct} \varepsilon_{ct} (h_{swi} - h_{a,i}), \quad (7.10)$$

where \dot{V}_a^{ct} denotes the airflow of cooling tower fans, h_{swi} the specific enthalpy of saturated air at temperature $T_{c,o}^{ads}$, and $h_{a,i}$ the specific enthalpy of air entering the cooling tower at a temperature T_{amb} . The cooling tower effectiveness ε_{ct} ,

$$\varepsilon_{ct} = \frac{1 - \exp \left[m \left(\frac{\rho_w \dot{V}_w^{ct}}{\rho_a \dot{V}_a^{ct}} \right)^n \left(\frac{\rho_a \dot{V}_a^{ct} C_{p,a}}{\rho_w \dot{V}_w^{ct} C_{p,w}} \right) \right]}{1 - \left(\frac{\rho_a \dot{V}_a^{ct} C_{p,a}}{\rho_w \dot{V}_w^{ct} C_{p,w}} \right) \cdot \exp \left[m \left(\frac{\rho_w \dot{V}_w^{ct}}{\rho_a \dot{V}_a^{ct}} \right)^n \left(\frac{\rho_a \dot{V}_a^{ct} C_{p,a}}{\rho_w \dot{V}_w^{ct} C_{p,w}} - 1 \right) \right]}, \quad (7.11)$$

where \dot{V}_w^{ct} is the waterflow in the cooling tower and equal to \dot{V}_w^{ads} , $m = 3$ and $n = 0.4$ are model constants obtained from [41, 42]. Various parameters used for modelling the thermal performance of the cooling tower are provided in Table 7.4.

The power consumed by the fans and pump used in the cooling tower air and water loops, respectively, are given by Eqs. (7.12) and (7.13),

$$\dot{P}_{fan}^{ct} = \frac{(K_a^{ct} \dot{V}_a^{ct})^2 \times \dot{V}_a^{ct}}{\eta_{fan}^{ct}}, \quad (7.12)$$

$$\dot{P}_{pump}^{ct} = \frac{(K_w^{ct} \dot{V}_w^{ct})^2 \times \dot{V}_w^{ct}}{\eta_{pump}^{ct}}, \quad (7.13)$$

where K_a^{ct} and K_w^{ct} are the pressure loss coefficients of airside and waterside, η_{fan}^{ct} and η_{pump}^{ct} are the efficiencies of the fans and pump (see Table 7.4).

Table 7.4: Parameters for the cooling tower and supporting systems used to extract heat from the adsorption chiller condenser [41-43].

Quantity and symbol	Value
Efficiency of cooling tower fans (η_{fan}^{ct})	0.8
Efficiency of cooling tower pump (η_{pump}^{ct})	0.8
Pressure loss coefficient for cooling tower air path (K_a^{ct})	0.12 Pa ^{0.5} m ⁻³ s
Pressure loss coefficient for cooling tower water loop (K_w^{ct})	1311 Pa ^{0.5} m ⁻³ s
Airflow of cooling tower fans (\dot{V}_a^{ct})	8000 CFM (3.78 m ³ s ⁻¹)
Waterflow in cooling tower loop ($\dot{V}_w^{ct} = \dot{V}_c^{ads}$)	50.2 GPM (3.167 × 10 ⁻³ m ³ s ⁻¹)
Dew point of air entering the cooling tower	20 °C
Relative humidity of air entering the cooling tower	0.88
Atmospheric pressure at the air inlet of the cooling tower	101.3 kPa
Cooling tower parameter (m)	3
Cooling tower parameter (n)	0.4

7.4.4 Energy analysis

The total power consumption for the two DC cooling solutions OS and NS, is obtained from Eqs. (7.14) and (7.15). For OS, the total cooling power consumption is from the liquid- and air-cooled racks, while for the NS, there is additional power consumption in the adsorption chiller \dot{P}_{ads} , cooling tower fans \dot{P}_{fan}^{ct} and pumps \dot{P}_{pump}^{ct} , i.e.,

$$\dot{P}_{OS} = \dot{P}'_{lc} + (\dot{P}'_{IRCU} + \dot{P}'_{vcr}), \quad (7.14)$$

$$\dot{P}_{NS} = \dot{P}''_{lc} + (\dot{P}''_{IRCU} + \dot{P}''_{vcr}) + (\dot{P}_{ads} + \dot{P}_{fan}^{ct} + \dot{P}_{pump}^{ct}), \quad (7.15)$$

where \dot{P}_{lc} denotes the power consumption that provides cooling for the liquid-cooled racks, \dot{P}_{IRCU} the consumption of IRCU fans obtained from the model of air-cooled racks, \dot{P}_{vcr} the consumption of VCR chiller, \dot{P}_{ads} power consumption of internal pumps of adsorption chiller. To realize the energy benefits of the NS over the existing infrastructure OS, the energy savings efficiency (ESE),

$$\text{ESE} = \left(\frac{\dot{P}_{OS} - \dot{P}_{NS}}{\dot{P}_{OS}} \right) \times 100. \quad (7.16)$$

7.4.5 Economic and environmental assessment

The liquid- and air-cooled racks are coupled through an adsorption chiller and associated subsystems, for which a cost-benefit analysis must be performed. The life cycle cost (LCC) of a DC cooling system has three components, (1) capital expenditure C_{cap} , (2) operational expenditure associated with electricity costs C_{op} , and (3) penalty costs due to CO₂ emissions C_{env} . The NS contains additional components compared to the OS, such as a silica gel-H₂O adsorption chiller, a pumping system for MT loop of adsorption chiller, cooling tower, and fans, which significantly increments C_{cap} . The total cost of ownership C_{total} [41],

$$C_{total} = \sum_l C_{cap}^l + C_{op} + C_{env}, \quad (7.17)$$

The first component, C_{cap}^l is the capital expenditure of the l^{th} equipment added to the existing system OS. C_{cap}^l depends on the sizing and capacity of the constituent equipment, for which the cost functions are listed in Table 7.5 [41, 44]. In Table 7.5, Q_{ads} and Q_{ct} are the cooling capacities of the adsorption chiller and cooling tower, respectively, \dot{P}_{fan} and

\dot{P}_{pump} are the rated power consumptions of fan and pump associated to the cooling tower, and η_{pump} is the efficiency of the pump.

Table 7.5: Capital expenditure (in CA\$) for the newly added equipment for waste heat recovery [41, 44].

Component	Capital cost function	CAPEX (LTC 30)	CAPEX (LTC 90)
Adsorption chiller	$C_{cap}^{ads} = 1590Q_e^{ads}$	≈ 46905 US\$	≈ 112706 US\$
Cooling tower	$C_{cap}^{ct} = 8815\ln(Q_{ct}) - 1749$	≈ 38845 Euro	≈ 44956 Euro
Fan	$C_{cap}^{fan} = (21.5644 + 20.1544\dot{P}_{fan}^{0.5})^2$	≈ 3000 Euro	≈ 3000 Euro
Pump	$C_{cap}^{pump} = 705.48\dot{P}_{pump}^{0.5} \left(1 + \frac{0.2}{1 - \eta_{pump}}\right)$	≈ 5000 US\$	≈ 5000 US\$
$\Delta C_{cap} = C_{cap}^{ads} + C_{cap}^{ct} + C_{cap}^{fan} + C_{cap}^{pump}$		≈ 124211 CA\$	≈ 216492 CA\$

The variable electricity tariff for Ontario, Canada, during different periods of the day is considered to calculate the OPEX component (see Table 7.6) [45].

Table 7.6: Variable electricity pricing in Ontario for on-peak, mid-peak, and off-peak hours [45].

Period	Price (CA\$/kWh)	Duration (h)
On-peak	0.176	6
Mid-peak	0.119	6
Off-peak	0.085	12

The annual OPEX for NS C_{op}^{NS} and OS C_{op}^{OS} based on the variable electricity tariff and other operational costs (including chemical and water costs for liquid-cooled DC),

$$\begin{aligned}
 C_{op}^{NS} = & \left[\dot{P}_{NS} \times (c_{on}t_{on} + c_{mid}t_{mid} + c_{off}t_{off}) \right] \\
 & + \{Q_{lc} \times (c_{chem} + c_{water}) \times 24\} \times 365,
 \end{aligned} \tag{7.18}$$

$$C_{op}^{OS} = \left[\left\{ \dot{P}_{OS} \times (c_{on}t_{on} + c_{mid}t_{mid} + c_{off}t_{off}) \right\} + \{Q_{lc} \times (c_{chem} + c_{water}) \times 24\} \right] \times 365, \quad (7.19)$$

where c_{on} , c_{off} , and c_{mid} are the electricity tariffs during on-peak t_{on} , off-peak t_{off} , and mid-peak t_{mid} periods, respectively (see Table 7.6). The values for the chemical c_{chem} and water c_{water} costs per kW of refrigeration-hour are provided in Table 7.1.

The third component, penalty cost associated with CO₂ emission C_{env} is considered within the total cost for the system NS and OS as follows [41].

$$C_{env}^{NS} = \frac{\mu_{CO_2} \times c_{CO_2} \times \dot{P}_{NS} \times 8760}{1000}, \quad (7.20)$$

$$C_{env}^{OS} = \frac{\mu_{CO_2} \times c_{CO_2} \times \dot{P}_{OS} \times 8760}{1000}, \quad (7.21)$$

where $\mu_{CO_2} = 0.968$ kg/kWh is the emission conversion factor of grid electricity utilized for powering the DC cooling systems and $c_{CO_2} = 112.5$ CA\$/ton of CO₂ is the penalty cost factor for electricity-associated CO₂ emission [41].

Another essential economic metric is the cost payback period CPP, which signifies the minimum time required to break even the capital expenditure from savings in operational expenditure. For this study, the reduction in OPEX and penalty costs is due to the profit obtained by utilizing the adsorption chiller. An increase in CAPEX ΔC_{cap} corresponds to the initial investment, as shown in Table 7.5. Thus, the payback period for the proposed waste heat recovery solution is,

$$CPP = \frac{\text{Investment}}{\text{Annual Savings}} = \frac{\Delta C_{cap}}{(C_{op}^{OS} + C_{env}^{OS}) - (C_{op}^{NS} + C_{env}^{NS})}. \quad (7.22)$$

7.5 Results and discussion

7.5.1 Parametric sensitivity analysis

The choice of systems parameters significantly impacts the ESE and annual running costs (including operational and CO₂ emission penalty costs). Four important parameters are considered here, (1) inlet water temperature to the liquid-cooled racks $T_{w,i}^{lc}$, (2) chilled water temperature supplied to the IRCUs of air-cooled racks $T_{w,i}^{IRCU}$, (3) total airflow of IRCUs \dot{V}_a^{IRCU} , and (4) re-cooling water temperature $T_{c,i}^{ads}$ supplied to the adsorption chiller condenser by the wet cooling tower. Out of the two choices of chillers in Tables 2 and 3, the results here are for a system with a TRANE CGA-120 VCR chiller and an InvenSor LTC30 e-plus adsorption chiller.

7.5.1.1 Influence of water inlet temperature of liquid-cooled infrastructure

In the NS, the inlet water to the liquid-cooled racks at $T_{w,i}^{lc}$ is heated with the server-generated heat and produces warmer water to drive the adsorption chiller. The heat added to the generator of the adsorption chiller depends on its exergy value or associated temperature. Table 7.7 and Figure 7.3 show the effect of raising the $T_{w,i}^{lc}$ on the ESE and annual operational expenditures for different (1) liquid-cooled rack heat dissipation Q_{lc} and (2) load sharing factor α .

Table 7.7: Parameters for the sensitivity analysis results, shown in Figure 7.3.

Parameter	Value
Inlet water temperature to liquid-cooled racks $T_{w,i}^{lc}$ ($^{\circ}\text{C}$)	40 to 65 (with 2.5 $^{\circ}\text{C}$ intervals)
Heat dissipation of liquid-cooled racks Q_{lc} (kW)	150 and 300
Waterflow of liquid-cooled racks \dot{V}_{lc} (GPM)	50 (for 150 kW), 100 (for 300 kW)
Chilled water setpoint temperature $T_{w,i}^{IRCU}$ ($^{\circ}\text{C}$)	15
Setpoint airflow of IRCUs \dot{V}_a^{IRCU} (CFM)	3290
Heat dissipation of air-cooled racks Q_{ac} (kW)	47
Re-cooling inlet water temperature $T_{c,i}^{ads}$ ($^{\circ}\text{C}$)	24
Ambient dry bulb temperature T_{amb} ($^{\circ}\text{C}$)	24
Load sharing factor α	0.2, 0.5, and 0.8
Allowable temperature of liquid-cooled CPUs ($^{\circ}\text{C}$)	85
Allowable cold chamber air temperature ($^{\circ}\text{C}$)	27

A high-level inspection of Figure 7.3a and b reveals that the ESE improves monotonically with increasing $T_{w,i}^{lc}$, due to improvement in COP_{ads} at elevated generator input temperatures [39]. However, raising $T_{w,i}^{lc} \geq 65$ $^{\circ}\text{C}$ violates the reliability constraint $\max[T_{CPU}^{lc}] \leq 85$ $^{\circ}\text{C}$ and thus damages the CPUs inside the liquid-cooled servers (see Eq. (7.2)). Lowering α produces a higher ESE due to the significantly reduced load on the VCR chiller (see Eq. (7.4)). An $\alpha = 0.5$ signifies that the VCR and the adsorption chillers handle equal fractions of the Q_{ac} . With $Q_{lc} = 150$ kW and $\dot{V}_{lc} = 50$ GPM, up to 16% ESE is achieved by raising $T_{w,i}^{lc}$, which offers considerable energy and economic savings. Comparison of Figure 7.3a and b shows that lowering Q_{lc} and \dot{V}_{lc} offers up to 6% additional ESE. However, \dot{V}_{lc} must be maintained above the sum of the design flowrates of liquid-cooled servers to satisfy $\max[T_{CPU}^{lc}] \leq 85$ $^{\circ}\text{C}$. The sum of annual running costs and the CO₂ emission penalty costs as a function of $T_{w,i}^{lc}$ are shown in Figure 7.3c and d. For all scenarios, the annual variable costs for the NS are lower than for the OS, which is more

pronounced as α is lowered. Hence, both $T_{w,i}^{lc}$ and α are essential to implement such a waste heat utilization strategy.

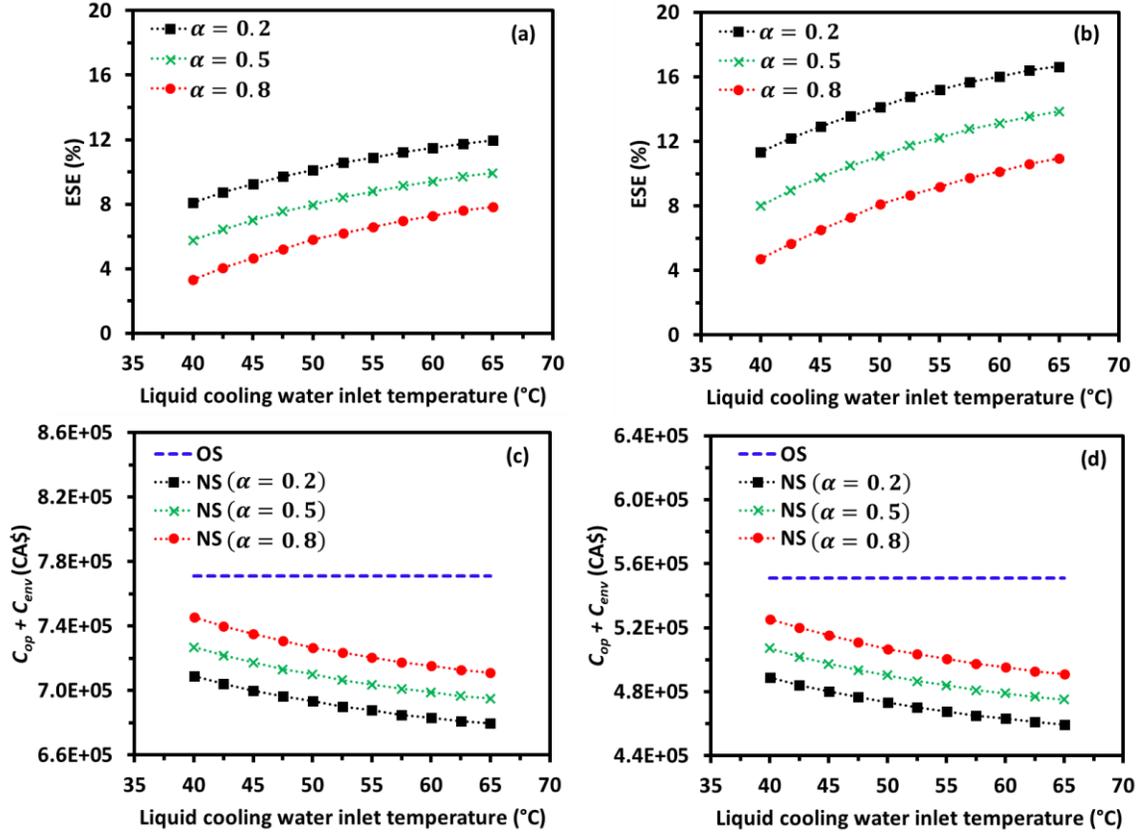


Figure 7.3: Energy savings efficiency (a, b) and annual costs (c, d) as a function of inlet water temperature of liquid-cooled racks $T_{w,i}^{lc}$ for $Q_{lc} = 300$ kW (left column) and $Q_{lc} = 150$ kW (right column). Each of the figures contains three different load sharing factors $\alpha = 0.2, 0.5,$ and 0.8 . Corresponding system parameters for the simulations are listed in Table 7.

7.5.1.2 Influence of chilled water temperature and airflow setpoints of air-cooled infrastructure

The chilled water produced at a specific temperature by the VCR and adsorption chillers is used to extract air-cooled rack-generated heat with the IRCUs. Parameters that control the cooling infrastructure are the chilled water temperature setpoint $T_{w,i}^{IRCU}$ and total airflow

setpoint \dot{V}_a^{IRCU} of the IRCUs. Figure 7.4 and Table 7.8 show the effect of increasing $T_{w,i}^{lc}$ on the ESE and annual operational expenditures for different Q_{lc} and α . In all simulations, $T_{w,i}^{IRCU}$ and \dot{V}_a^{IRCU} are adjusted so that $\max[T_{j,k}^f] \leq 27$ °C to ensure thermally reliable operation of the air-cooled servers.

Table 7.8: Parameters for the sensitivity analysis results, shown in Figure 7.4.

Parameter	Value
Inlet water temperature to liquid-cooled racks $T_{w,i}^{lc}$ (°C)	55
Heat dissipation of liquid-cooled racks Q_{lc} (kW)	150 and 300
Waterflow of liquid-cooled racks \dot{V}_{lc} (GPM)	50 (for 150 kW), 100 (for 300 kW)
Chilled water setpoint temperature $T_{w,i}^{IRCU}$ (°C)	8 to 20 (with 2 °C intervals)
Setpoint airflow of IRCUs \dot{V}_a^{IRCU} (CFM)	2900 to 3600
Heat dissipation of air-cooled racks Q_{ac} (kW)	47
Re-cooling inlet water temperature $T_{c,i}^{ads}$ (°C)	24
Ambient dry bulb temperature T_{amb} (°C)	24
Load sharing factor α	0.2, 0.5, and 0.8
Allowable temperature of liquid-cooled CPUs (°C)	85
Allowable cold chamber air temperature (°C)	27

To couple the two cooling parameters $T_{w,i}^{IRCU}$ and \dot{V}_a^{IRCU} a dimensionless product $\theta \times \beta$ is used [33, 46], where

$$\theta = \frac{T_{w,i}^{IRCU}}{T_{amb}} \text{ and } \beta = \frac{\dot{V}_a^{IRCU}}{125 \times Q_{ac}}. \quad (7.23)$$

Here, θ denotes the dimensionless chilled water temperature and β the dimensionless airflow of IRCUs. The denominator of β is based on the ASHRAE airflow guidelines for air-cooled DCs, which specify 125 CFM per unit kW heat load [47]. Intuitively, if θ is increased, the corresponding dimensionless airflow β must also be increased to satisfy the constraint $\max[T_{j,k}^f] \leq 27$ °C.

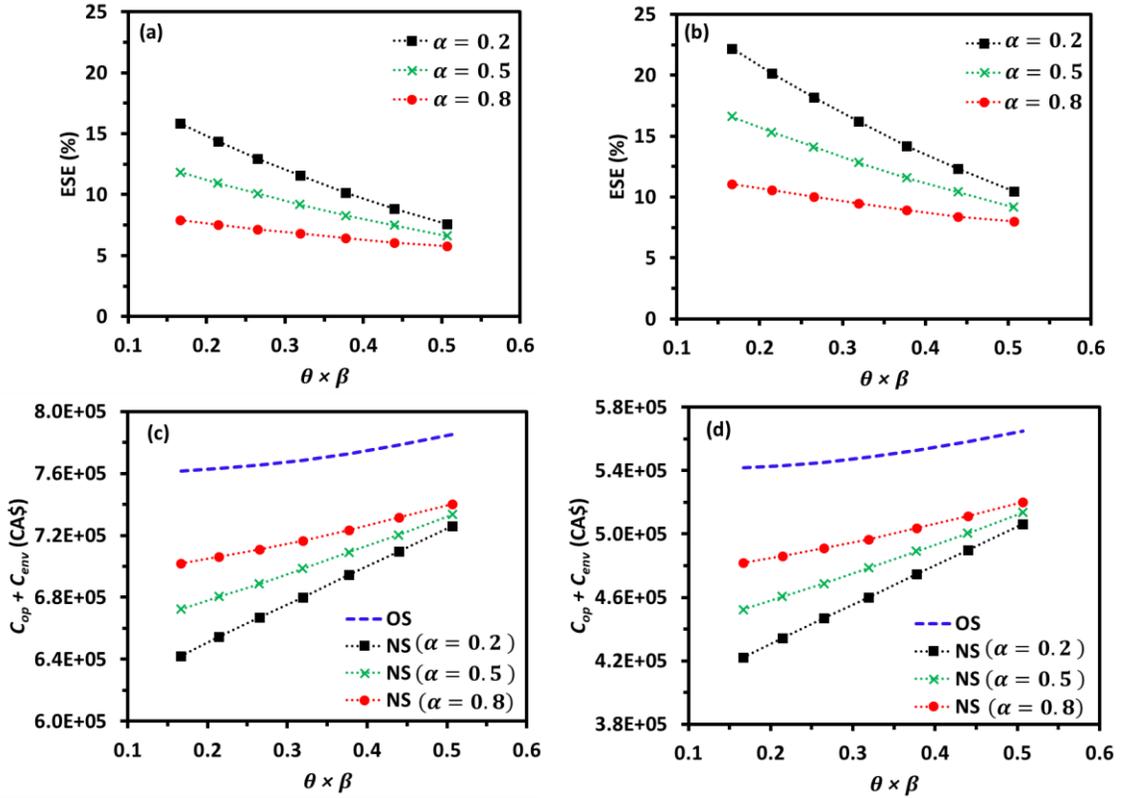


Figure 7.4: Energy savings efficiency (a, b) and annual costs (c, d) as a function of the dimensionless product $\theta \times \beta$ for $Q_{lc} = 300$ kW (left column) and $Q_{lc} = 150$ kW (right column). Each of the figures contains three different load sharing factors $\alpha = 0.2, 0.5$, and 0.8 . Corresponding system parameters for the simulations are listed in Table 8.

Figure 7.4 shows that increasing the dimensionless product $\theta \times \beta$ monotonically decreases the ESE (Figure 7.4a and b), raising the annual sum of operational and penalty costs (Figure 7.4c and d). At higher values of α , up to 22% of ESE is achieved by regulating $\theta \times \beta$, due to higher fraction of cooling provided by the adsorption chiller. Two reasons for ESE improvement at lower $\theta \times \beta$ are that (1) lowering θ increases the temperature difference between the evaporator and adsorption chiller generator, offering better heat transfer efficiency and (2) lowering β reduces fan power consumption of the IRC units. Figure 7.4c and d show that the savings in annual variable costs for the NS are lower than

for the OS, where savings are more realizable with lower α . Therefore, the dimensionless cooling parameter $\theta \times \beta$ and load sharing factor α can be used to optimize the thermodynamic and economic performance of the NS.

7.5.1.3 Influence of re-cooling circuit water inlet temperature to adsorption chiller

Another parameter controlling the energy efficiency of the waste heat recovery system is the inlet water temperature $T_{c,i}^{ads}$ of the adsorption chiller re-cooling circuit supplied by the wet cooling tower. To ensure energy balance across the control volume, the condenser must reject the added heat to the generator and the evaporator of the adsorption chiller (see Eq. (7.8)). This is done by water supplied by the wet cooling tower at $T_{c,i}^{ads}$. Table 7.9 and Figure 7.5 show the effect of increasing $T_{c,i}^{ads}$ on the ESE and annual operational expenditures for different (1) liquid-cooled rack heat dissipation Q_{lc} and (2) load sharing factors α .

Table 7.9: Parameters for the sensitivity analysis results, shown in Figure 7.5.

Parameter	Value
Inlet water temperature to liquid-cooled racks $T_{w,i}^{lc}$ (°C)	55
Heat dissipation of liquid-cooled racks Q_{lc} (kW)	150 and 300
Waterflow of liquid-cooled racks \dot{V}_{lc} (GPM)	50 (for 150 kW), 100 (for 300 kW)
Chilled water setpoint temperature $T_{w,i}^{IRCU}$ (°C)	14
Setpoint airflow of IRCUs \dot{V}_a^{IRCU} (CFM)	3232
Heat dissipation of air-cooled racks Q_{ac} (kW)	47
Re-cooling inlet water temperature $T_{c,i}^{ads}$ (°C)	22 to 36 (with 2 °C interval)
Ambient dry bulb temperature T_{amb} (°C)	24
Load sharing factor α	0.2, 0.5, and 0.8
Allowable temperature of liquid-cooled CPUs (°C)	85
Allowable cold chamber air temperature (°C)	27

Figure 7.5a and b show that raising $T_{c,i}^{ads}$ does not influence ESE significantly for all Q_{lc} and α . At $\alpha = 0.2$ and $\alpha = 0.8$, ESE increases and decreases, respectively, changing ESE at most by 2%. For $\alpha = 0.2$, a significant fraction of the heat load of air-cooled racks is handled by the adsorption chiller, whereas for $\alpha = 0.8$, this is reversed. When the cooling load on the adsorption chiller is lower, a higher $T_{c,i}^{ads}$ can lower ESE. Figure 7.5a and b corroborate Figure 7.5c and d, where a higher ESE corresponds to annual variable costs savings ($C_{op} + C_{env}$) for the NS over the OS.

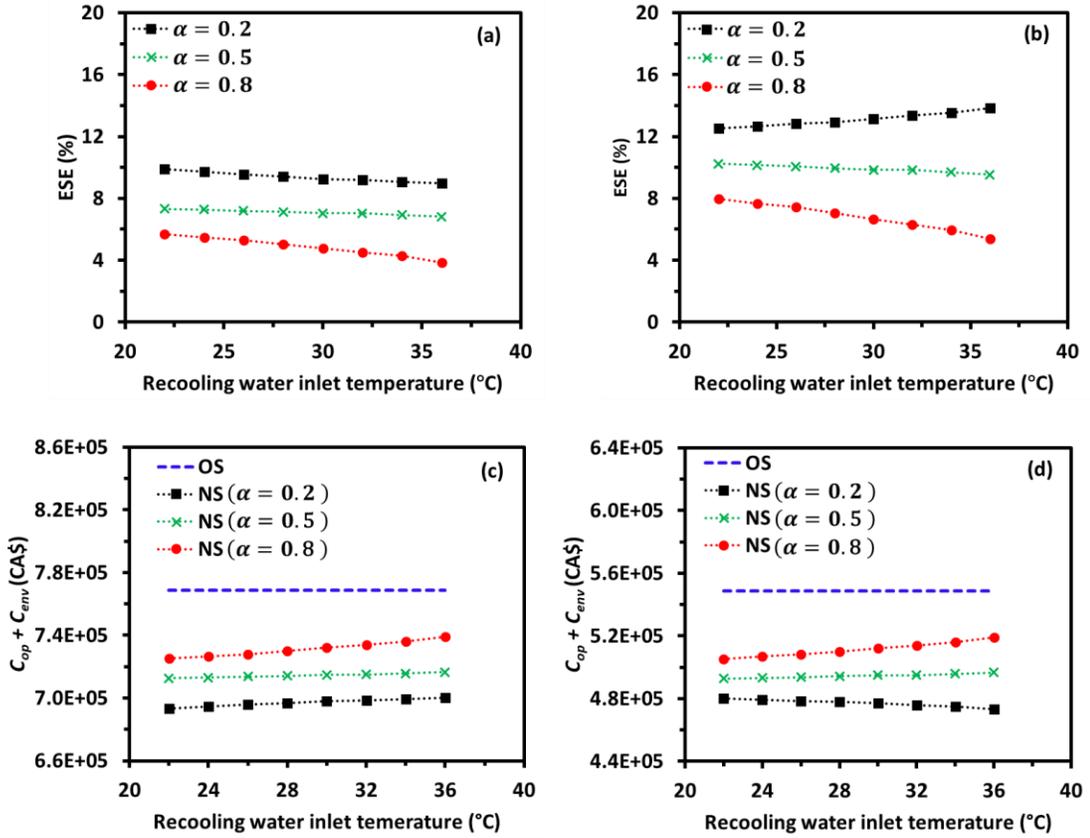


Figure 7.5: Energy savings efficiency (a, b) and annual costs (c, d) as a function of re-cooling water inlet temperature of adsorption chiller condenser $T_{c,i}^{ads}$ for (a) $Q_{lc} = 300$ kW (left column) and $Q_{lc} = 150$ kW (right column). Each of the figures contains three different load sharing factors $\alpha = 0.2, 0.5$, and 0.8 . Corresponding system parameters for the simulations are listed in Table 9.

7.5.2 Economic and environmental performance evaluation

A waste heat recovery system must be energy-efficient and environmentally sustainable yet provide the shortest return on investment (ROI). For the NS, the initial investment in the waste heat recovery solution is the capital expenditure associated with the equipment listed in Table 7.5, and the return occurs through the annual accumulation of operational cost savings. The metric to quantify the ROI is the CPP, shown in Eq. (7.22). Both operational and environmental costs due to grid electricity-associated CO₂ emissions are included in the CPP. A lower CPP value is desirable, which ensures that the NS will make a profit within the lowest possible operational span. The metric quantifying the environmental sustainability of the NS is the annual reduction in tons of grid electricity associated CO₂ emission defined as $\Delta m_{CO_2} = m_{CO_2}^{OS} - m_{CO_2}^{NS}$ (see Eqs. (20) and (21)).

Different scenarios from the prior parametric investigations are selected for Δm_{CO_2} and CPP to investigate the degradation (or augmentation) of these metrics, as shown in Table 7.10. Scenarios 1 to 9 correspond to variations in $T_{w,i}^{lc}$, scenarios 10 to 18 examine the influence of $\theta \times \beta$, while Scenarios 19 to 21 raise $T_{c,i}^{ads}$. For all cases, increasing α monotonically increases the CPP. This is because, at higher α , the cooling load of the air-cooled racks handled by the adsorption chiller is lower than for the VCR chiller, resulting in higher values of \dot{P}_{lc} and \dot{P}_{VCR} . For lower α a significant fraction of these two components of power consumption is reduced, resulting in a shorter CPP and higher Δm_{CO_2} . The lowest possible CPP (379 days) and highest Δm_{CO_2} (44.8 tons) is obtained for Scenario 10 from the highest ESE corresponding to that scenario. Therefore, for a given $T_{w,i}^{lc}$ and $T_{c,i}^{ads}$, it is

instructive to regulate the dimensionless air cooling parameter $\theta \times \beta$ to ensure thermodynamic and economic efficiency of the NS.

Table 7.10: CPP and Δm_{CO_2} for a wide range of scenarios with InvenSor LTC30 e-plus (29.5 kW) adsorption chiller retrofit. The constant parameters for these scenarios are $\dot{V}_{lc} = 100$ GPM, $Q_{lc} = 300$ kW, $Q_{ac} = 47$ kW, and $T_{amb} = 24$ °C.

Scenario	$T_{w,i}^{lc}$ (°C)	$T_{w,i}^{IRCU}$ (°C)	\dot{V}_a^{IRCU} (CFM)	$T_{c,i}^{ads}$ (°C)	α	Δm_{CO_2} (tons)	CPP (days)
1	50	15	3290	24	0.2	16.6	585
2	50	15	3290	24	0.5	22.7	747
3	50	15	3290	24	0.8	29.0	1023
4	60	15	3290	24	0.2	20.8	517
5	60	15	3290	24	0.5	26.9	631
6	60	15	3290	24	0.8	32.9	816
7	65	15	3290	24	0.2	22.5	497
8	65	15	3290	24	0.5	28.4	598
9	65	15	3290	24	0.8	34.2	757
10	55	8	2950	24	0.2	44.8	379
11	55	8	2950	24	0.5	33.5	507
12	55	8	2950	24	0.8	22.4	759
13	55	14	3232	24	0.2	30.6	513
14	55	14	3232	24	0.5	23.7	648
15	55	14	3232	24	0.8	16.9	875
16	55	20	3590	24	0.2	13.4	770
17	55	20	3590	24	0.5	10.6	880
18	55	20	3590	24	0.8	8.1	1011
19	55	14	3232	34	0.2	25.9	656
20	55	14	3232	34	0.5	19.8	858
21	55	14	3232	34	0.8	12.2	1392

7.5.3 Scalability of the proposed solution for data centers with larger capacity

The scalability of NS to a larger air cooling capacity is investigated. The VCR and adsorption chillers are the TRANE CGAF-C25 and InvenSor LTC90 e-plus, where each

chiller has ≈ 88 kW design cooling capacity (see Table 7.2 and Table 7.3). For the liquid-cooled racks $Q_{lc} = 300$ kW heat load, and the heat loads of the air-cooled racks are assumed to be double that in the previous sections, i.e., $Q_{lc} = 94$ kW. Considering $Q_{ads} \approx 88$ kW and $Q_{ct} \approx 200$ kW, the capital expenditure for the larger cooling system is therefore $\Delta C_{cap} \approx 216492$ CA\$ (see Table 7.5). Since air-cooled DCs already pre-exist with a VCR chiller, this is not considered in the ΔC_{cap} . The aim is to investigate the ESE, Δm_{CO_2} , and CPP for the 88 kW adsorption chiller retrofit as a function of $T_{w,i}^{lc}$, $\theta \times \beta$, and $T_{c,i}^{ads}$.

Table 7.11 shows the results for 21 different scenarios where the $T_{w,i}^{lc}$, $T_{w,i}^{IRCU}$, \dot{V}_a^{IRCU} , $T_{c,i}^{ads}$, and α are varied for the DC equipped with 88 kW VCR and adsorption chillers. Again, Scenario 10 provides the lowest CPP (285 days), highest ESE (22.6%) and Δm_{CO_2} (≈ 104 tons). This improvement in energy efficiency of the 88 kW system compared to the 30 kW system corresponds to higher waste heat capturing from the exhaust water stream of the liquid-cooled racks. Hence, installing an adsorption chiller with a cooling capacity close to the maximum heat dissipation of the liquid-cooled rack improves the thermodynamic and economic performance of the NS.

Table 7.11: ESE, CPP, and Δm_{CO_2} for a wide range of scenarios with InvenSor LTC90 e-plus (88.5 kW) adsorption chiller retrofit. The constant parameters for these scenarios are $\dot{V}_{lc} = 100$ GPM, $Q_{lc} = 300$ kW, $Q_{ac} = 94$ kW, and $T_{amb} = 24$ °C.

Scenario	$T_{w,i}^{lc}$ (°C)	$T_{w,i}^{IRCU}$ (°C)	\dot{V}_a^{IRCU} (CFM)	$T_{c,i}^{ads}$ (°C)	α	ESE (%)	Δm_{CO_2} (tons)	CPP (days)
1	50	15	6580	24	0.2	11.4	51.7	573
2	50	15	6580	24	0.5	8.7	39.4	752
3	50	15	6580	24	0.8	6.0	27.2	1088
4	60	15	6580	24	0.2	14.2	64.3	461
5	60	15	6580	24	0.5	11.6	52.5	564
6	60	15	6580	24	0.8	8.9	40.3	735
7	65	15	6580	24	0.2	15.2	68.7	431
8	65	15	6580	24	0.5	12.6	57.2	517
9	65	15	6580	24	0.8	10.0	45.3	654
10	55	8	5900	24	0.2	22.6	103.9	285
11	55	8	5900	24	0.5	17.7	81.4	364
12	55	8	5900	24	0.8	12.8	59.0	502
13	55	14	6464	24	0.2	17.2	72.0	367
14	55	14	6464	24	0.5	14.3	58.6	441
15	55	14	6464	24	0.8	11.3	44.5	559
16	55	20	7180	24	0.2	12.0	34.2	508
17	55	20	7180	24	0.5	10.9	28.7	561
18	55	20	7180	24	0.8	9.9	23.9	616
19	55	14	6464	34	0.2	13.0	60.9	486
20	55	14	6464	34	0.5	9.7	45.4	652
21	55	14	6464	34	0.8	6.5	30.5	972

7.5.4 Degradation of computing efficiency due to waste heat recovery

We explore the effect of system parameters on the computing performance per unit power consumption in the NS. Following the LINPACK performance test for servers, the Floating Point Operations per second (FLOPs) per watt is considered for this purpose. The FLOPs/W metric depends on several factors, such as (1) operating voltage, (2) clock

frequency, and (3) chip temperature. Since raising the liquid cooling inlet temperature $T_{w,i}^{lc}$ increases CPU temperature (see Eq. (7.2)), the computing performance is affected. GFLOPs/W as a function of T_{CPU}^{lc} for liquid-cooled servers with the Intel Xeon E5 family of processors is [35],

$$\text{GFLOPs/W} = -0.0069T_{CPU}^{lc} + 3.0118. \quad (7.24)$$

Figure 7.6 shows the tradeoffs between ESE and normalized GFLOPs/W as a function of $T_{w,i}^{lc}$. Intuitively, raising $T_{w,i}^{lc}$ increases ESE over a wide range of α . However, this degrades GFLOP/W due to the increase in CPU temperature. At higher values of T_{CPU}^{lc} , the computing chip undergoes thermal throttling so that the leakage current of the chip increases. As a result, the CPU provides lower computing performance (GFLOPs) at higher power consumption. Hence the DCs with such a waste heat extraction mechanism must continuously monitor both energy and computing efficiency to fulfill multiple objectives.

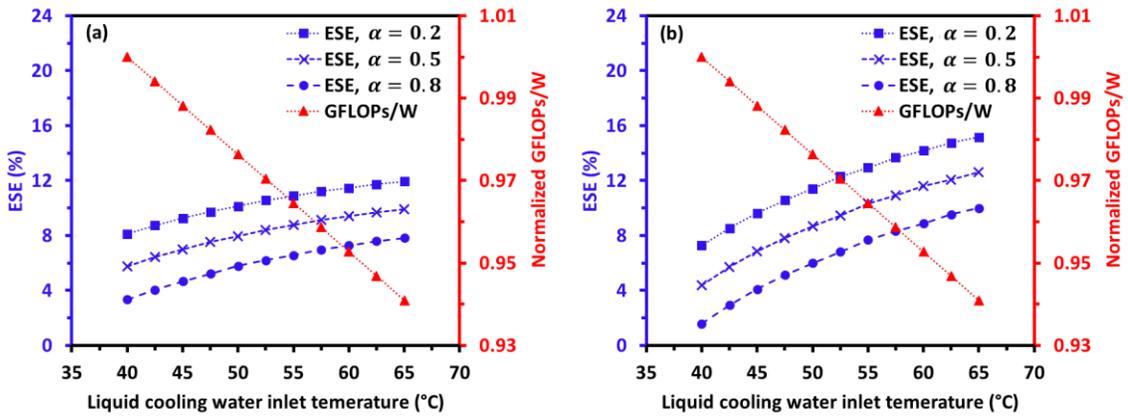


Figure 7.6: Tradeoffs between computing performance per watt (GFLOPs/W) and energy savings efficiency (ESE) as a function of inlet water temperature of liquid-cooled racks $T_{w,i}^{lc}$ for InvenSor (a) LTC30 e-plus and (b) LTC90 e-plus adsorption chiller-based waste heat recovery.

7.6 Conclusion

A strategy to utilize DC waste heat is developed for heterogeneous infrastructures that contain both air- and liquid-cooled racks. The architecture utilizes waste heat internally within a DC with minimal retrofit and ensures energy efficiency. Thermodynamic, economic, and environmental performance of the heat recovery solution is assessed while discussing potential retrofit in existing DCs. The findings include:

- The waste heat recovery solution can offer up to 23% ESE through system parameter choices.
- Higher water inlet temperatures of liquid-cooled racks improve ESE by improving adsorption chiller COP.
- Increasing the dimensionless cooling parameter of air-cooled racks $\theta \times \beta$ lowers ESE and results in higher operational expenditure.
- A scalability analysis shows that retrofitting a 300 kW heat-generating liquid-cooled DC room with 29.5 and 88.5 kW adsorption chillers significantly shortens payback periods to 379 and 285 days, respectively.
- The new strategy enables annual electricity savings up to 96 MWh, which translates to an annual reduction of 104 tons of CO₂ equivalent greenhouse gas emissions.
- By decreasing the load sharing factor α from 0.8 to 0.2, the ESE improves up to 10% due to savings in the power consumption of the VCR chiller.

- Selecting an adsorption chiller with a cooling capacity near the expected average heat load of liquid-cooled racks improves ESE, reduces operational expenditures, and results in a shorter payback period.
- However, improving ESE by raising the liquid-cooling inlet temperature from 40°C to 65°C degrades the GFLOPs/W of liquid-cooled racks by 6%.

7.7 References

- [1] E. Masanet, A. Shehabi, N. Lei, S. Smith, and J. Koomey, "Recalibrating global data center energy-use estimates," *Science*, vol. 367, pp. 984-986, 2020.
- [2] W. Strielkowski, I. Firsova, I. Lukashenko, J. Raudeliūnienė, and M. Tvaronavičienė, "Effective management of energy consumption during the COVID-19 pandemic: The role of ICT solutions," *Energies*, vol. 14, p. 893, 2021.
- [3] R. Gupta, S. Asgari, H. Moazamigoodarzi, S. Pal, and I. K. Puri, "Cooling architecture selection for air-cooled Data Centers by minimizing exergy destruction," *Energy*, vol. 201, p. 117625, 2020.
- [4] H. Moazamigoodarzi, P. J. Tsai, S. Pal, S. Ghosh, and I. K. Puri, "Influence of cooling architecture on data center power consumption," *Energy*, vol. 183, pp. 525-535, 2019.
- [5] A. Almoli, A. Thompson, N. Kapur, J. Summers, H. Thompson, and G. Hannah, "Computational fluid dynamic investigation of liquid rack cooling in data centres," *Applied energy*, vol. 89, pp. 150-155, 2012.

- [6] H. Moazamigoodarzi, R. Gupta, S. Pal, P. J. Tsai, S. Ghosh, and I. K. Puri, "Modeling temperature distribution and power consumption in IT server enclosures with row-based cooling architectures," *Applied Energy*, vol. 261, p. 114355, 2020.
- [7] H. Moazamigoodarzi, S. Pal, S. Ghosh, and I. K. Puri, "Real-time temperature predictions in it server enclosures," *International Journal of Heat and Mass Transfer*, vol. 127, pp. 890-900, 2018.
- [8] H. Moazamigoodarzi, S. Pal, D. Down, M. Esmalifalak, and I. K. Puri, "Performance of a rack mountable cooling unit in an IT server enclosure," *Thermal Science and Engineering Progress*, vol. 17, p. 100395, 2020.
- [9] S. Asgari, S. MirhoseiniNejad, H. Moazamigoodarzi, R. Gupta, R. Zheng, and I. K. Puri, "A gray-box model for real-time transient temperature predictions in data centers," *Applied Thermal Engineering*, vol. 185, p. 116319, 2021.
- [10] S. Zimmermann, I. Meijer, M. K. Tiwari, S. Paredes, B. Michel, and D. Poulikakos, "Aquasar: A hot water cooled data center with direct energy reuse," *Energy*, vol. 43, pp. 237-245, 2012.
- [11] Lenovo, "ThinkSystem SD650 Direct Water Cooled Server, Product Guide," ed, 2018.
- [12] DELL, "PowerEdge C6525, Technical Guide," ed, 2019.
- [13] K. Ebrahimi, G. F. Jones, and 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*, vol. 31, pp. 622-638, 2014.

- [14] A. H. Khalaj and S. K. Halgamuge, "A Review on efficient thermal management of air-and liquid-cooled data centers: From chip to the cooling system," *Applied energy*, vol. 205, pp. 1165-1188, 2017.
- [15] B. B. Kanbur, C. Wu, S. Fan, and F. Duan, "System-level experimental investigations of the direct immersion cooling data center units with thermodynamic and thermoeconomic assessments," *Energy*, vol. 217, p. 119373, 2021.
- [16] C. Donnelly. (2020, April 22). *Datacentre liquid cooling: What needs to happen for it to become commonplace in colocation?* Available: www.computerweekly.com/feature/Datacentre-liquid-cooling-What-needs-to-happen-for-it-to-become-commonplace-in-colocation
- [17] A. Patrizio. (2020, April 22). *Ready to dive into water cooling?* Available: www.datacenterdynamics.com/en/analysis/ready-dive-water-cooling/
- [18] E. Oró, P. Taddeo, and J. Salom, "Waste heat recovery from urban air cooled data centres to increase energy efficiency of district heating networks," *Sustainable cities and society*, vol. 45, pp. 522-542, 2019.
- [19] P. Huang, B. Copertaro, X. Zhang, J. Shen, I. Löfgren, M. Rönnelid, *et al.*, "A review of data centers as prosumers in district energy systems: Renewable energy integration and waste heat reuse for district heating," *Applied Energy*, vol. 258, p. 114109, 2020.

- [20] T. Brunschwiler, G. I. Meijer, S. Paredes, W. Escher, and B. Michel, "Direct waste heat utilization from liquid-cooled supercomputers," in *International Heat Transfer Conference*, 2010, pp. 429-440.
- [21] M. Romero, H. Hasselqvist, and G. Svensson, "Supercomputers keeping people warm in the winter," in *2nd International Conference on ICT for Sustainability (ICTS), AUG 24-27, 2014, Stockholm, SWEDEN*, 2014, pp. 324-332.
- [22] Z. He, T. Ding, Y. Liu, and Z. Li, "Analysis of a district heating system using waste heat in a distributed cooling data center," *Applied Thermal Engineering*, vol. 141, pp. 1131-1140, 2018.
- [23] J. Yu, Y. Jiang, and Y. Yan, "A simulation study on heat recovery of data center: A case study in Harbin, China," *Renewable energy*, vol. 130, pp. 154-173, 2019.
- [24] M. Deymi-Dashtebayaz and S. Valipour-Namanlo, "Thermoeconomic and environmental feasibility of waste heat recovery of a data center using air source heat pump," *Journal of Cleaner Production*, vol. 219, pp. 117-126, 2019.
- [25] A. R. Murphy and A. S. Fung, "Techno-economic study of an energy sharing network comprised of a data centre and multi-unit residential buildings for cold climate," *Energy and Buildings*, vol. 186, pp. 261-275, 2019.
- [26] K. Ebrahimi, G. F. Jones, and A. S. Fleischer, "The viability of ultra low temperature waste heat recovery using organic Rankine cycle in dual loop data center applications," *Applied Thermal Engineering*, vol. 126, pp. 393-406, 2017.
- [27] S. Araya, G. F. Jones, and A. S. Fleischer, "Organic Rankine Cycle as a Waste Heat Recovery System for Data Centers: Design and Construction of a Prototype," in

2018 17th IEEE Intersociety Conference on Thermal and Thermomechanical Phenomena in Electronic Systems (ITherm), 2018, pp. 850-858.

- [28] M. R. J. Al-Tameemi, Y. Liang, and Z. Yu, "Combined ORC-HP thermodynamic cycles for DC cooling and waste heat recovery for central heating," *Energy Procedia*, vol. 158, pp. 2046-2051, 2019.
- [29] J. B. Marcinichen, J. A. Olivier, and J. R. Thome, "On-chip two-phase cooling of datacenters: Cooling system and energy recovery evaluation," *Applied Thermal Engineering*, vol. 41, pp. 36-51, 2012.
- [30] A. Haywood, J. Sherbeck, P. Phelan, G. Varsamopoulos, and S. K. Gupta, "Thermodynamic feasibility of harvesting data center waste heat to drive an absorption chiller," *Energy Conversion and Management*, vol. 58, pp. 26-34, 2012.
- [31] K. Ebrahimi, G. F. Jones, and A. S. Fleischer, "Thermo-economic analysis of steady state waste heat recovery in data centers using absorption refrigeration," *Applied Energy*, vol. 139, pp. 384-397, 2015.
- [32] H. Chen, Y.-h. Peng, and Y.-l. Wang, "Thermodynamic analysis of hybrid cooling system integrated with waste heat reusing and peak load shifting for data center," *Energy Conversion and Management*, vol. 183, pp. 427-439, 2019.
- [33] R. Gupta, H. Moazamigoodarzi, S. MirhoseiniNejad, D. G. Down, and I. K. Puri, "Workload management for air-cooled data centers: An energy and exergy based approach," *Energy*, vol. 209, p. 118485, 2020.

- [34] A. Carbó, E. Oró, J. Salom, M. Canuto, M. Macías, and J. Guitart, "Experimental and numerical analysis for potential heat reuse in liquid cooled data centres," *Energy conversion and management*, vol. 112, pp. 135-145, 2016.
- [35] A. A. Moskovsky, E. A. Druzhinin, A. B. Shmelev, V. V. Mironov, and A. Semin, "Server Level Liquid Cooling: Do Higher System Temperatures Improve Energy Efficiency?," *Supercomputing frontiers and innovations*, vol. 3, pp. 67-74, 2016.
- [36] M. Sahini, U. Chowdhury, A. Siddarth, T. Pradip, D. Agonafer, R. Zeigham, *et al.*, "Comparative study of high ambient inlet temperature effects on the performance of air vs. liquid cooled IT equipment," in *2017 16th IEEE Intersociety Conference on Thermal and Thermomechanical Phenomena in Electronic Systems (ITherm)*, 2017, pp. 544-550.
- [37] J. Gordon and K. C. Ng, "Predictive and diagnostic aspects of a universal thermodynamic model for chillers," *International Journal of Heat and Mass Transfer*, vol. 38, pp. 807-818, 1995.
- [38] TRANE, "Air-cooled liquid chillers, 10 to 60 Tons," ed, 2004.
- [39] InvenSor, "LTC 30 e plus, Product datasheet," ed, 2019.
- [40] InvenSor, "LTC 90 e plus, Product datasheet," ed, 2019.
- [41] A. H. Khalaj, T. Scherer, and S. K. Halgamuge, "Energy, environmental and economical saving potential of data centers with various economizers across Australia," *Applied energy*, vol. 183, pp. 1528-1549, 2016.
- [42] A. H. Khalaj, K. Abdulla, and S. K. Halgamuge, "Towards the stand-alone operation of data centers with free cooling and optimally sized hybrid renewable

power generation and energy storage," *Renewable and Sustainable Energy Reviews*, vol. 93, pp. 451-472, 2018.

- [43] M. Iyengar and R. Schmidt, "Analytical modeling for thermodynamic characterization of data center cooling systems," *Journal of Electronic Packaging*, vol. 131, 2009.
- [44] A. Wyrwa and Y.-K. Chen, "Techno-Economic Analysis of Adsorption-Based DH Driven Cooling System," in *IOP Conference Series: Earth and Environmental Science*, 2019, p. 012128.
- [45] O. E. Board. (April 22). *Historical electricity rates*. Available: www.oeb.ca/rates-and-your-bill/electricity-rates/historical-electricity-rates
- [46] R. Gupta, D. G. Down, and I. K. Puri, "Energy and Exergy-Aware Workload Assignment for Air-Cooled Data Centers," in *Proceedings of the Twelfth ACM International Conference on Future Energy Systems (e-Energy'21)*, Virtual Event, Italy, 2021.
- [47] K. Brown, W. Torell, and V. Avelar, "Choosing the optimal data center power density," *Schneider White, Boston, MA*, p. 2, 2014.

8 Conclusions

Thermal and workload management has always been an important research problem in the DC industry. Therefore, the dissertation develops energy and exergy-based analyses informed by computational fluid dynamics simulations to choose optimized air distribution schemes intelligently. The equipment-level exergy loss analysis also quantifies cooling inefficiencies at a granular level through exergy destruction. This helps identify the loss of available cooling energy for a range of cooling architectures, operating parameters, and strategies by which optimal airflow can be achieved. It was observed that the legacy cooling architectures with underfloor air delivery and overhead air delivery schemes are up to 23% less exergy-efficient compared to the modular cooling geometries with in-row cooling and rack-mountable cooling. It was also shown that the DC exergy destruction is a function of Peclet number and dimensionless dead state temperature ratio. The compression chiller used in such air-cooled DCs contributes as high as 50% of the overall exergy loss. This loss can be reduced by two means (1) recovering waste heat at the hot air exhaust stream of the chiller condenser, or (2) replacing the compression chiller with a thermally driven chiller that consumes way less energy to produce chilled water.

Subsequently, for the exergy-efficient geometries, i.e., in-row cooling and rack-mountable cooling, workload and cooling management strategies having been developed by fulfilling multiple objectives such as energy, exergy, and computing efficiencies. Low-complexity spatiotemporal physics-based zonal thermal models informed the optimization problems solved for these two DCs. Trade-offs across different objectives, workload

parameters, cooling parameters, and server thermal reliability guidelines were thoroughly investigated for single-rack and multi-rack DCs. The multi-objective Pareto optimal fronts between different objective functions help realize potential savings from different DC operation perspectives.

In the final stage, the exergy destruction component of the vapor compression chiller was reduced through an internal waste heat utilization strategy for data centers with both air and liquid –cooled racks. Waste heat was recovered from the exhaust hot water stream of the liquid-cooled racks using an adsorption chiller, which provided on-site production of chilled water. This chilled water source could either fully/partially reduce the cooling load on the vapor compression chiller. Techno-economic analysis showed that retrofitting a data center with such waste heat recovery infrastructure can have a payback period as low as 284 days and offer an annual reduction of 104 tons of CO₂ equivalent greenhouse gas emissions. However, while implementing such a waste heat recovery solution, specific operating conditions can degrade computing performance per watt at the cost of improving waste heat recovery efficiency.

9 Future directions

- The proposed single- and multi-objective optimization problems addressing energy, exergy, computing efficiencies can be implemented in the real-time DC control system through economic-model predictive control (MPC) and adaptive predictive control (APC) strategies.
- The exergy destruction due to recirculation and bypass in various air-cooled DCs can be controlled in real-time by workload-aware fan control systems.
- Since incorporating the server consolidation will drastically change the optimization and thermal framework, developing a robust consolidation strategy addressing these objectives remains an open problem.
- As an alternative strategy to that proposed in Chapter 7, the exergy loss suffered by the vapor compression chillers of air-cooled DCs can be reduced by recovering waste heat from the airside of the chiller condenser. For this purpose, exergo-economic analysis is required parallel to DC thermal analysis to evaluate overall system performance.
- For air-cooled DCs equipped with a waste heat extraction subsystem, a waste heat utilization-aware workload and cooling management framework will be essential.

10 Appendices

10.1 Performance data of the vapor compression chillers for various condenser air inlet temperatures

In chapters 4,5,6, and 7, the thermodynamic analysis of the entire DC cooling system requires an accurate model to simulate the chiller's off-design performance that produces the chilled water needed for the air-handler units (UFAD, OHAD, IRC, or RMCU). For this purpose, a widely used Gordon-Ng model [1] was selected due to its low computational requirement while accurately capturing the part-load chiller efficiency. For all the cases, we use a vapor-compression chiller with an ambient air-cooled condenser. The model takes the following form.

$$y = A_1x_1 + A_2x_2 + A_3x_3, \quad (9.1)$$

$$y = \left[\left(1 + \frac{1}{COP} \right) \frac{T_{c,w}}{T_{amb}} \right], \quad x_1 = \frac{T_{c,w}}{Q_{ch}}, \quad x_2 = \frac{T_{amb} - T_{c,w}}{Q_{ch}T_{amb}}, \quad x_3 = \frac{\left(1 + \frac{1}{COP} \right) Q_{ch}}{T_{amb}}. \quad (9.2)$$

where \dot{Q}_{ch} is the heat load on the chiller, $COP = \dot{Q}_{ch}/\dot{P}_{ch}$ is the coefficient of performance, \dot{P}_{ch} is the total power consumption of the chiller (i.e., the sum of the compressor, condenser blowers, and internal circulating pump of the chiller), $T_{c,w}$ is the chilled water temperature produced by the chiller evaporator, and T_{amb} is the ambient air temperature entering the chiller condenser. In this case A_1 , A_2 , and A_3 are performance constants required to be obtained based on a specific chiller datasheet. In Chapter 4, a TRANE CGAF C25 (60 Hz) chiller with 25 Ton designed refrigeration capacity was used. The manufacturer-provided performance data of the chiller is shown in Figure 10.1 [2]. It is important to note that all

the temperatures are in K in this model, whereas \dot{Q}_{ch} and \dot{P}_{ch} are in kW. By performing non-linear regression using the data presented in Figure 10.1, following model constants are obtained $A_1 = 0.146$, $A_2 = 42.097$, and $A_3 = 1.861$. Other important parameters required for the thermodynamic models are listed in Table 10.1.

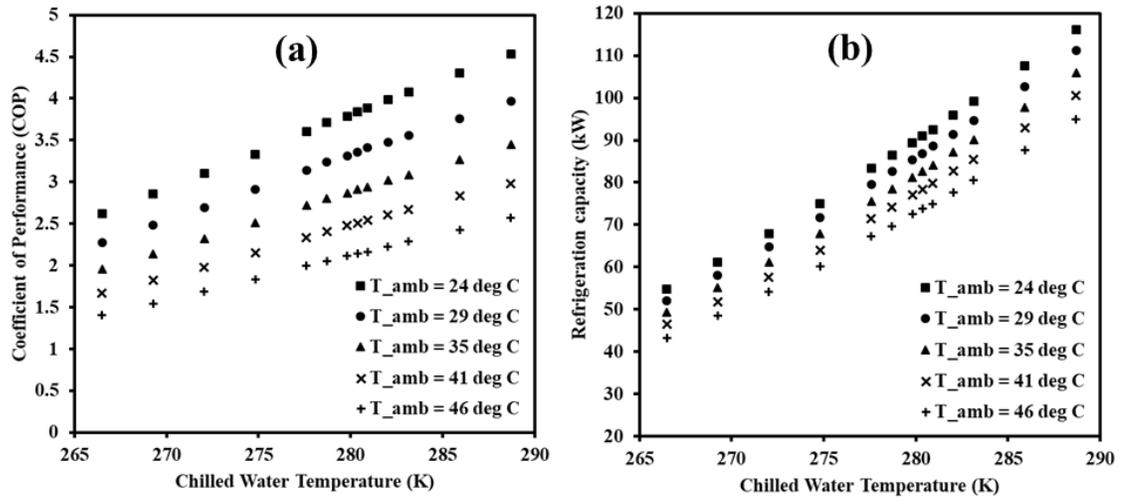


Figure 10.1: Chiller (a) coefficient of performance and (b) refrigeration capacity as a function of chilled water temperature for different condenser air inlet temperatures. The data is obtained from the datasheet of TRANE CGAF C25 (60 Hz) chiller with 25 Ton designed refrigeration capacity.

Table 10.1: Additional specifications for TRANE CGAF C25 (60 Hz) chiller.

Quantity	Details
Compressor type	Scroll
Nominal refrigeration capacity	25 ton
Minimum water flow	30 GPM
Maximum water flow	90 GPM
Total airflow of condenser fans	21650 CFM

In Chapters 5 and 6, the compression chiller model TRANE CGA 120 (60 Hz) with 10 Ton designed refrigeration capacity was used to produce chilled water for the IRC and

RMCU heat exchangers. Figure 10.2 shows the performance data for the chiller obtained from the manufacturer [2]. By performing non-linear regression using the data presented in Figure 10.2, the model constants are $A_1 = 0.054$, $A_2 = 17.267$, and $A_3 = 4.952$.

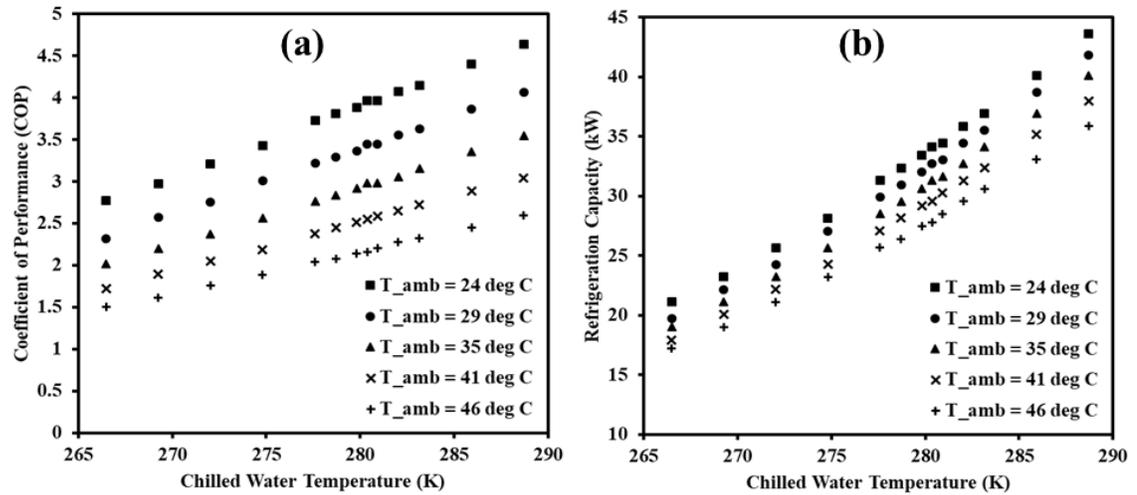


Figure 10.2: (a) Chiller coefficient of performance and (b) refrigeration capacity as a function of chilled water temperature for different condenser air inlet temperatures. The performance data is obtained from the datasheet of TRANE CGA 120 (60 Hz) chiller with 10 Ton designed refrigeration capacity.

Table 10.2: Additional specification for TRANE CGA 120 (60 Hz) chiller.

Quantity	Details
Compressor type	Scroll
Nominal refrigeration capacity	10 ton
Minimum water flow	12 GPM
Maximum water flow	36 GPM
Total airflow of condenser fans	8120 CFM

10.2 Detailed modeling strategy of a single server zone for the zonal temperature prediction approach

The transient characteristic of a single server zone used in the zonal temperature prediction model is elaborated in this section. Since intricate component-level modeling of all the electronic components inside the server chassis (such as CPUs, DIMMs, server fans, power supply modules, PCIe cards, HDDs, and SSDs) will significantly increase the execution time of the model, a zonal approach is adopted for representing the thermo-fluid the server chassis [3-8]. The aim is to capture transient physical phenomenon with minimal computational time (within a few seconds). Figure 10.3 shows a representative schematic of salient temperatures, airflows, and energy interactions across the outer control volume of a server zone with its adjacent front (cold) and back (hot) chamber zones.

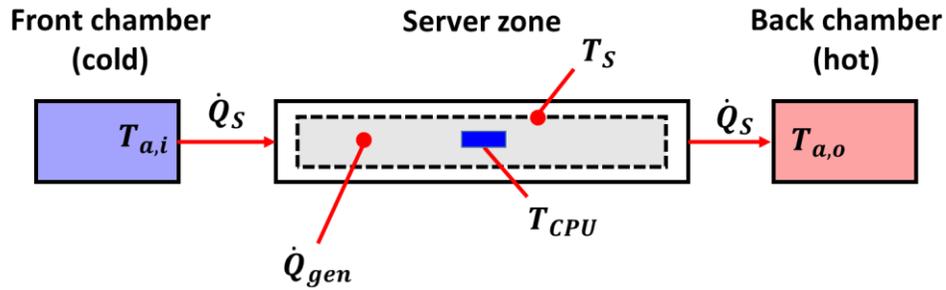


Figure 10.3: Representative schematic showing salient temperatures, airflows, and energy interactions for a single server zone interacting with front (cold) and back (hot) chambers. The total volume of the server zone is identical to that of a 1U server chassis.

By imposing a transient energy balance across the outer control volume of a server zone, one can write [3, 9, 10],

$$X_a \frac{dT_{a,i}}{dt} + X_s \frac{dT_s}{dt} = hA(T_s - T_{a,i}) + \chi \dot{P}_s, \quad (9.3)$$

where $T_{a,i}$ and $T_{a,o}$ the inlet and exhaust air temperatures, T_s the mean effective boundary temperature of the server control volume, T_{CPU} the average surface temperature of computing chips, hA is the product of convective heat transfer coefficient and area, $X_a = \rho_a C_{p,a} V_a$ and $X_s = \rho_s C_s V_s$ are the thermal masses of air and solid material of the server, \dot{P}_s is the electrical power consumption of the server, and χ is an empirical constant expressed as the ratio between heat transfer by the cold air stream to the power consumption of the server so that $0 \leq \chi \leq 1$. Furthermore, the \dot{P}_s can be expressed in terms of percentage of server utilization so that $\dot{P}_s = C_1 + C_2 U$. However, several unknown parameters such as hA , χ , C_1 , and C_2 are required to be measured to obtain a closure form Eq. 9.3 to predict T_s . Equation 9.3 can be rewritten by replacing hA with R_{th} so that the thermal resistance of a server $R_{th} = 1/hA$. The product $X_s \times R_{th}$ is known as the time constant of a server τ_s , which indicates the transient evolution of temperature profiles [3, 9, 10].

$$\frac{dT_s}{dt} + \frac{X_a}{X_s} \left[\frac{dT_{a,i}}{dt} \right] = \frac{1}{X_s \times R_{th}} [T_s - T_{a,i}] + \frac{\chi}{X_s} [C_1 + C_2 U], \quad (9.4)$$

The values of X_a and X_s are constant for a given physical dimension of a server chassis expressed in terms of ‘U’ (1U = 4.36 cm) found in the literature [3]. However, the values for R_{th} , χ , C_1 , and C_2 are very specific to the manufacturer, model, and internal components of a server. In all the cases investigated, an HP ProLiant DL360 G5 server was used, which is a 1U air-cooled server [11]. Table 10.3 shows the results of U , \dot{P}_s , $T_{a,i}$, $T_{a,o}$, T_{CPU} , and \dot{Q}_s , obtained from in-house performed experiments for the aforementioned server model.

Table 10.3: Benchmark experiments for the HP ProLiant DL360 G5 server to find unknown parameters in the thermal model.

U (%)	$T_{a,i}$ (°C)	\dot{P}_s (W)	$T_{a,o}$ (°C)	T_{CPU} (°C)	\dot{Q}_s (m ³ s ⁻¹)
0	19.1	226.83	28.6	35	0.0184
10	19.1	235.44	29.3	39	0.0184
20	19.1	251.76	30.1	43	0.0184
30	19.1	269.04	31	47	0.0184
40	19.1	286.83	31.8	52	0.0184
50	19.1	300.22	32.5	55	0.0184
60	19.1	317.10	33.3	59	0.0184
70	19.1	332.45	34	62	0.0184
80	19.1	350.37	34.9	66	0.0184
90	19.1	363.65	35.6	71	0.0184
100	19.1	373.58	35.9	73	0.0184
0	22.0	221.09	31	37	0.0184
10	22.0	236.13	32	41	0.0184
20	22.0	254.07	32.8	46	0.0184
30	22.0	269.96	33.7	51	0.0184
40	22.0	287.45	34.4	53	0.0184
50	22.0	302.63	35.4	58	0.0184
60	22.0	317.87	36.1	62	0.0184
70	22.0	334.11	37	66	0.0184
80	22.0	351.37	38	71	0.0184
90	22.0	365.49	38.5	73.5	0.0184
100	22.0	374.50	39	76	0.0184
0	24.6	219.77	34.1	40	0.0184
10	24.6	233.85	34.9	43.5	0.0184
20	24.6	250.36	35.7	48	0.0184
30	24.6	266.47	36.4	54	0.0184
40	24.6	283.95	37.2	58	0.0184
50	24.6	300.51	38.4	62	0.0184
60	24.6	317.23	39.1	66	0.0184
70	24.6	333.66	39.8	70	0.0184
80	24.6	350.17	40.6	73	0.0184
90	24.6	365.51	41.1	75	0.0184
100	24.6	374.08	41.4	77	0.0184
0	26.6	221.47	35.3	41.5	0.0198

10	26.6	236.93	35.9	45	0.0198
20	26.6	253.26	36.6	49	0.0198
30	26.6	270.17	37.2	56	0.0198
40	26.6	286.76	38	59	0.0198
50	26.6	303.51	38.7	62	0.0198
60	26.6	319.79	39.4	65	0.0198
70	26.6	336.23	40.2	69	0.0198
80	26.6	350.42	40.8	72	0.0198
90	26.6	371.30	41.4	75	0.0198
100	26.6	370.87	41.6	77	0.0198
0	28.5	220.07	36.3	42.5	0.0222
10	28.5	235.69	36.8	46	0.0222
20	28.5	252.91	37.4	50	0.0222
30	28.5	271.47	38.1	54	0.0222
40	28.5	283.56	38.6	57	0.0222
50	28.5	301.30	39.4	60	0.0222
60	28.5	317.83	40.2	64	0.0222
70	28.5	334.09	40.9	69	0.0222
80	28.5	350.18	41.6	71	0.0222
90	28.5	366.10	42	74	0.0222
100	28.5	376.36	42.4	77	0.0222
0	30.3	223.57	37.8	44	0.02454
10	30.3	237.31	38.3	47	0.02454
20	30.3	255.41	38.6	50	0.02454
30	30.3	271.17	39.3	54	0.02454
40	30.3	287.30	39.9	58	0.02454
50	30.3	304.88	40.6	61	0.02454
60	30.3	319.67	41.2	65	0.02454
70	30.3	335.99	41.8	69	0.02454
80	30.3	352.72	42.1	72	0.02454
90	30.3	375.58	42.5	74	0.02454
100	30.3	377.92	42.8	76	0.02454
0	32.7	224.39	39.3	45	0.0269
10	32.7	238.45	39.9	48	0.0269
20	32.7	256.36	40.5	53	0.0269
30	32.7	273.67	40.9	57	0.0269
40	32.7	290.51	41.5	59	0.0269

50	32.7	307.76	41.8	62	0.0269
60	32.7	322.95	42.5	66	0.0269
70	32.7	339.95	43	69	0.0269
80	32.7	356.03	43.5	73	0.0269
90	32.7	373.48	44	75	0.0269
100	32.7	379.15	44.5	76	0.0269

From Table 10.3, several important parameters are plotted in Figure 10.4 to understand the influence of our primary variables U and $T_{a,i}$ on the model parameters. Figure 10.4a shows the variation of server power as functions of U and $T_{a,i}$. It is evident that the server power remains nearly invariant if the $T_{a,i}$ increases. Therefore, the prior hypothesis $\dot{P}_s(U) = C_1 + C_2U$ holds true. Within the range of $0 \leq U(\%) \leq 100$ the constants obtained by curve fitting are $C_1 = 224.9$ and $C_2 = 1.6$.

Next, we plot the thermal resistance expressed as $R_{th} = \left\{ T_{CPU} - \left(\frac{T_{a,i} + T_{a,0}}{2} \right) \right\} / \dot{P}_s$ as functions of U and $T_{a,i}$ in Figure 10.4b. For this case, a non-linear dependence on the input variables $U(\%)$ and $T_{a,i}(\text{°C})$ are observed and is represented by the equation $R_{th}(U, T_{a,i}) = B_1 + B_2U + B_3T_{a,i} + B_4U^2 + B_5UT_{a,i}$. The constants are obtained by non-linear regression resulting in $B_1 = 0.064$, $B_2 = 1.49 \times 10^{-3}$, $B_3 = -6.88 \times 10^{-4}$, $B_4 = -5.59 \times 10^{-6}$, and $B_5 = -1.04 \times 10^{-5}$.

Similarly, the CPU temperature T_{CPU} is plotted as functions of U and $T_{a,i}$ in Figure 10.4b to reveals its underlying variations. For this case as well, we assume the functional form $T_{CPU}(U, T_{a,i}) = D_1 + D_2U + D_3T_{a,i} + D_4U^2 + D_5UT_{a,i}$ and find $D_1 = 20.21$, $D_2 = 0.59$, $D_3 = 0.76$, $D_4 = -8.14 \times 10^{-4}$, and $D_5 = -5.6 \times 10^{-3}$ using non-linear regression.

Finally, in Figure 10.4d, the dimensionless quantity $\chi = \{\rho_a \dot{Q}_s C_{p,a} (T_{a,o} - T_{a,i})\} / \dot{P}_s$, which represents the ratio of server-generated heat to the input electric power. It is observed from Figure 10.4d that the majority of the χ values fall within a tight band of $0.91 \leq \chi \leq 1$. Therefore, a mean value $\chi = 0.955$ is assumed in the entire thermal model. The aforementioned equations obtained from a regression of the experimental data provide closure to the thermal resistance-capacitance-based model of a single server zone (see Eq. 10.4) while accurately representing its behavior.

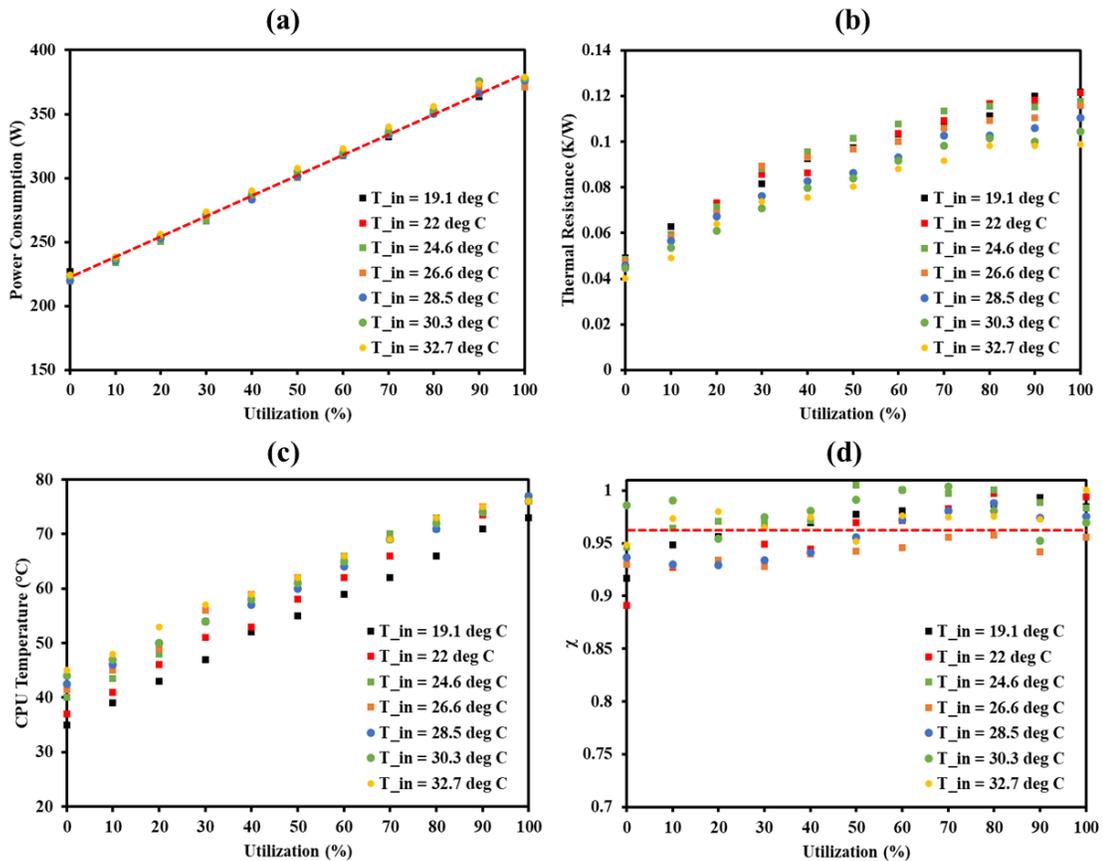


Figure 10.4: Experimentally measured (a) power consumption, (b) thermal resistance, (c) CPU temperature, and (d) χ as functions of a percentage of utilization supplied with different inlet air temperatures for HP ProLiant DL360 G5 server.

Figure 10.5 shows the influence of server air inlet temperature on the airflow of the server obtained from Table 10.3. It is observed that for $T_{a,i} \leq 27^\circ\text{C}$, the server airflow $\dot{Q}_s \approx 0.0184 \text{ m}^3 \text{ s}^{-1}$ remains nearly invariant. Since the optimization problems solved in Chapters 5 and 6 de facto considers a capping on the cold chamber temperature, i.e., $T_{a,i} \leq 27^\circ\text{C}$ by following the ASHRAE guideline [12], including the influence of $T_{a,i}$ on \dot{Q}_s is not necessary for this purpose. However, we note that for accurately examining data center cooling cycle faults where $T_{a,i}$ can rise well above 27°C , considering Figure 10.5 becomes utmost essential.

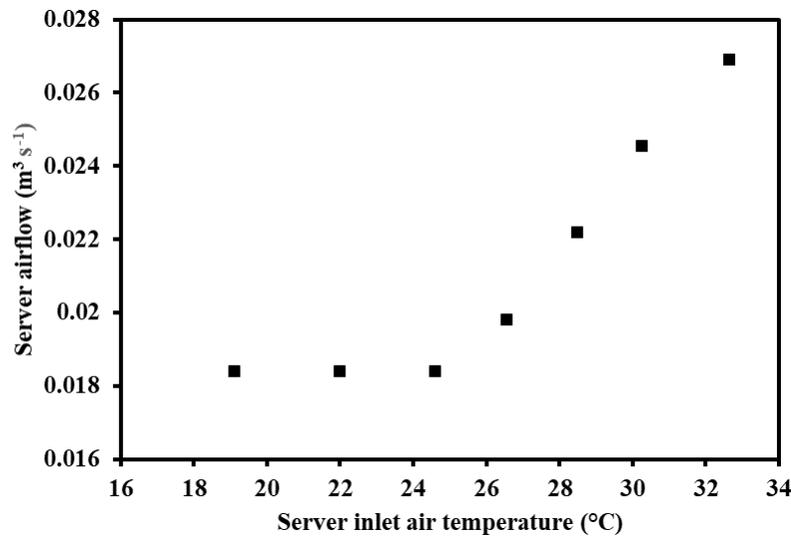


Figure 10.5: Total airflow of server fans as a function of inlet air temperature for HP ProLiant DL360 G5 server.

10.3 References

- [1] J. M. Gordon, K. Ng, and H. Chua, "Optimizing chiller operation based on finite-time thermodynamics: universal modeling and experimental confirmation," *International Journal of Refrigeration*, vol. 20, pp. 191-200, 1997.

- [2] TRANE, "Air-cooled liquid chillers, 10 to 60 Tons," ed, 2004.
- [3] J. W. VanGilder, Z. Pardey, C. Healey, and X. Zhang, "A compact server model for transient data center simulations," *ASHRAE Transactions*, vol. 119, pp. 358-370, 2013.
- [4] K. Fouladi, A. P. Wemhoff, L. Silva-Llanca, K. Abbasi, and A. Ortega, "Optimization of data center cooling efficiency using reduced order flow modeling within a flow network modeling approach," *Applied Thermal Engineering*, vol. 124, pp. 929-939, 2017.
- [5] S. Asgari, S. MirhoseiniNejad, H. Moazamigoodarzi, R. Gupta, R. Zheng, and I. K. Puri, "A gray-box model for real-time transient temperature predictions in data centers," *Applied Thermal Engineering*, vol. 185, p. 116319, 2021.
- [6] S. Asgari, H. Moazamigoodarzi, P. J. Tsai, S. Pal, R. Zheng, G. Badawy, *et al.*, "Hybrid surrogate model for online temperature and pressure predictions in data centers," *Future Generation Computer Systems*, vol. 114, pp. 531-547, 2021.
- [7] H. Moazamigoodarzi, R. Gupta, S. Pal, P. J. Tsai, S. Ghosh, and I. K. Puri, "Modeling temperature distribution and power consumption in IT server enclosures with row-based cooling architectures," *Applied Energy*, vol. 261, p. 114355, 2020.
- [8] H. Moazamigoodarzi, S. Pal, S. Ghosh, and I. K. Puri, "Real-time temperature predictions in it server enclosures," *International Journal of Heat and Mass Transfer*, vol. 127, pp. 890-900, 2018.
- [9] Z. M. Pardey and J. W. VanGilder, "Further exploration of a compact transient server model," in *Fourteenth Intersociety Conference on Thermal and*

Thermomechanical Phenomena in Electronic Systems (ITherm), 2014, pp. 1322-1329.

- [10] Z. M. Pardey, D. W. Demetriou, H. S. Erden, J. W. VanGilder, H. E. Khalifa, and R. R. Schmidt, "Proposal for standard compact server model for transient data center simulations," in *ASHRAE Conference-Papers*, 2015, pp. 413-421.
- [11] H. Packard, "Quickspecs: HP ProLiant DL360 Generation 5 (G5)," in *Product datasheet*, 27, ed, 2007, pp. 1-43.
- [12] R. Steinbrecher and R. Schmidt, "Data center environments," *ASHRAE J*, vol. 42, p. e49, 2011.