Holistic Study of Thermal Management in Direct Liquid Cooled Data Centres: from the Chip to the Environment

Mustafa Alaa Kadhim Kadhim

Submitted in accordance with the requirements for the degree of Doctor of Philosophy

The University of Leeds

Institute of Thermo-Fluids

School of Mechanical Engineering

April, 2018

Papers contributing to this thesis

- Kadhim, M.A., Al-Anii, Y.T., Kapur, N., Summers, J.L. and Thompson, H.M., 2017, March. Performance of a mixed mode air handling unit for direct liquid-cooled servers. In Thermal Measurement, Modeling & Management Symposium (SEMI-THERM), 2017 33rd (pp. 172-178). IEEE.

- Kadhim, M.A, Kapur, N., Summers, J.L. and Thompson, H.M., 2017, Sepember. Effective heat transfer from microchips to the external environment. In UK Heat Transfer Conference, 2017 15[th] (PP 213). ISBN: 9781908549327.

- Kadhim, M.A, Kapur, N., Summers, J.L. and Thompson, H.M., 2018. Rack level study of hybrid liquid/air cooled servers: The impact of flow distribution and pumping configuration on the CPUs temperature. Accepted journal article to the journal of Heat Transfer Engineering (Taylor & Francis).

- Kadhim, M.A, Thompson, H.M., Summers, J.L. and Kapur, N, 2018. An experimental and theoretical investigation of droplet evaporation on heated hydrophilic and hydrophobic surfaces. Submitted journal article to Langmuir.

The candidate has collaborated with the co-authors as shown above. The supervisors have provided valuable review, contribution and guidance to the candidate.

Acknowledgements

Abstract

The IT (Information Technology) infrastructure power consumption constitutes a large portion of global electricity consumption and a large proportion of this energy is to maintain an acceptable thermal environment for the IT equipment. Therefore, it is important to understand and improve the thermal and energy management of data centres for lower cost and higher sustainability. Toward this goal, Direct Contact Liquid Cooled (DCLC) servers, where liquid loop heat exchangers are attached to the CPU, were proposed to study the use of chiller-less energy efficient data centre. Thirty Sun Fire V20z servers in a data centre rack have their CPUs water cooled with the remaining components air cooled, together with a rear door heat exchanger to capture this air heat flow. The heat generated by the servers is ultimately transferred to the environment using an Air Handling Unit (AHU). The AHU was fitted with a water spray system to increase the heat transfer capacity.

The designed DCLC system was tested and characterised in terms of power consumption and thermal performance. The design successfully provided stable inlet coolant temperature ($\pm 1°C$) to the IT despite the variation in the IT workload and environmental conditions. Activating the spray reduced the thermal resistance of the AHU heat exchanger (HE) by 50%. However, the power consumption and pressure drop across the HE was increased.

The flow distribution and the coolant pumping configurations of centralised (where the coolant is pumped by two central pumps connected in series) and distributed (where small pumps inside the servers are activated) was investigated. The EPANET software was used to analyse the flow and showed that the servers in the top of the rack receive a higher flow rate (by approximately 30%) than the servers in the bottom of the rack. This resulted in a variation in the CPU temperatures of different servers. Optimisation analysis proposed increasing the manifolds size to improve the flow rate and reduce the flow maldistribution. In the distributed pumping case, the CPUs temperature showed to be $2°C$ higher compared with the central pumping case for the high IT workload.

The rack inlet temperature was tested in the range of the ASHRAE W4 envelope in terms of CPU temperatures, power consumption and computational efficiency. Increasing the coolant inlet temperature resulted in high energy saving in the AHU, while the rack energy consumption increases marginally in idle operation and considerably more in high IT workloads. This results in an improvement in the energy effectiveness of 17% but a deterioration in the computational efficiency of 4%.

Finally, a parallel study was carried out to investigate the droplet evaporation over heated surfaces which ultimately be used in studying sprays in the AHU or in direct on chip cooling via evaporation. A novel experimental design was proposed to track the lifetime of any droplet size that span the surface tension to gravity-dominated regimes. A theoretical model was also proposed to predict the droplet lifetime based on the initial contact angle, contact radius and the receding contact angle. The model predicted the droplet evaporation over hydrophobic surfaces with good accuracy of an error less than 4% while under estimated the evaporation with hydrophilic surfaces.

Table of Contents

List of Tables

List of Figures

xix

Abbreviations

Acronyms

AC	Alternative Current
AHU	Air Handling Unit
AMD	Advanced Micro Devices
ASHRAE	American Society of Heating Refrigeration and Air Conditioning Engineers
CCA	Constant Contact Angle
CCR	Constant Contact Radius
CCS	Control and Collect System
CDU	Coolant Distributed Unit
CHx40	Liquid/liquid heat exchanger
CI	Capture Index
CMOS	Complementary Metal Oxide Semiconductor Transistor
COP	Coefficient of Performance
COPG	Grand Coefficient of Performance
CP	Central Pumping
CPU	Central Processing Unit
CRAC/CRAH	Computer Air Conditioning Unit
CUE	Carbon Usage Effectiveness
DC	Direct Currents
DCEPM	Data Centre Energy Performance Productivity
DCiE	Data Centre Infrastructure Efficiency
DCLC	Direct Contact Liquid Cooling
DDCP	Deep Drilled Cold Plate

DHCP	Dynamic Host Configuration Protocol
DoE	Design of Expert
DP	Distributed Pumping
DWPE	Data Centre Work Load Power Efficiency
ERE	Energy Reuse Effectiveness
ETS	Energy-to-Solution
FEP	Fluorinated Ethylene Propylene
FNM	Flow Network Modelling
FSC	Fan Speed Control
FTCP	Formed Tube Cold Plate
FVER	Fixed to Variable Energy Ratio
HDD	Hard Disk Drive
HE	Heat Exchanger
HPC	High-Performance Computers
HVAC	Heating Ventilation and Air Condoning
ICT	Information and communications technology
IPMI	Intelligent Platform Management Interface
IT	Information Technology
ITUE	IT-Total Usage Effectiveness
JIT	Just in Time Compilers
JVM	Java Virtual Machine
LCD	Liquid Crystal Display
LMTIM	Liquid Metal Thermal Interface Material
MACA	Media Access Control Address
MCCP	Machined Channel Cold Plate

PCB	Printed Circuit Board
pDWPE	Partial Data Centre Work Load Power Efficiency
PFCP	Pocketed Folded Fin Cold Plate
PI	Proportional Integrated
PID	Proportional Integrated Derivative
pPUE	Partial Power Usage Effectiveness
PSU	Power Supply Unit
PTDaemon	Power and Temperature Daemon
TCUE	Total Cooling Usage Effectiveness
PUE	Power Usage Effectiveness
PV	Process Variable
RAM	Random Access Memory
RCI	Rack Cooling Index
RDHx	Rear Door Heat Exchanger
REE	Relative Energy Efficiency
RHI	Return Heat Index
RTI	Return Temperature Index
RTPF	Round Tube Plain Fin
SCHx	Sidecar Heat Exchanger
SE	Spray Efficiency
SEPM	Server Energy Performance Productivity
SHI	Supply Heat Index
SP	Service Processer
SPEC	Standard Performance Evaluation Corporation
sPUE	System Power Usage Effectiveness

SSH	Secure Shell
SSL	Secure Sockets Layer
SSR	Stick-Slip Receding Angle Mode
SUT	Server Under Test
TDP	Thermal Design Power
TGI	The Green Index
TIM	Thermal Interface Material
TUE	Total Power Usage Effectiveness
UPS	Uninterrupted Power Supply Units
VR	Voltage Regulator
WPE	Workload Power Efficiency

Nomenclature

Terms	Units	Description
Q_c	W	Cooling capacity
E_{IT}	W	Power consumed by the IT
E_d	W	Direct power embedded in water
E_n	W	Indirect power embedded in water
U	mm	Rack unit $= 44.45$mm
h	$\dfrac{W}{m^2}.K$	Convection heat transfer coefficient
f	-	Friction factor
Q	l/s	Flow rate
P	Pa	Pressure
RH	%	Relative humidity
$\dot{V}$	m³/s	Air volume flow rate
A_d	m²	The cross sectional area
$\overline{U}$	m/s	Average air velocity
u_i	m/s	Measured velocity in each point in the duct
R	$°C/kW$	Thermal resistance

C_p	$\frac{J}{kg.K}$	Specific heat
$\dot{m}$	kg/s	Mass flow rate
X	-	Humidity ratio
$P_{cooling}$	W	Power consumed by cooling
P_{AHU}	W	Power consumed by the AHU
P_{CHx}	W	Power consumed by CHx40
$P_{process\ pump}$	W	Power consumed by the process water pump
P_{fan}	W	Power consumed by the AHU fan
$P_{spray\ pump}$	W	Power consumed by the spray pump
$P_{control\ unit}$	W	Power consumed by the control unit of the AHU
$P_{cpu\ pumps}$	W	Power consumed by the servers module pumps
$P_{servers\ fans}$	W	Power consumed by the small fans inside the servers
U	%	Utilisation workload
F_s	rpm	Fan speed
B_o	-	Bond number
D	m^2/s	Diffusion coefficient
c	kg/m^2	Vapour concentration
h_{fg}	kJ/kg	Latent heat of evaporation

Greek symbols

Terms	Units	Description
ω	-	Water usage energy
ρ	kg/m^3	Density
g	m/s^2	Gravitational acceleration
γ	N	Surface tension
L	m	Characteristics length
θ	degree	Contact angle
θ_r	degree	Receding contact angle

Chapter 1 Introduction

1.1 General overview

Data centres are computing infrastructures that contain information and communications technology (ICT or IT) equipment as shown in figure 1.1. The IT equipment (the function of the data centre) is used for storing and processing data and transmitting and receiving across the data centre boundary. Besides the IT equipment, data centres also house network routers, network connections, power supplies, data storages, switches, fire suppressions, high level security systems and cooling facilities [1, 2].

Many of our daily activities rely on data centres, this includes banks, government installations, educational institutions, online searching, communications and transportation. Hence, many applications worldwide are adversely affected if any malfunction happens in data centres; proving very costly to end-users and data centre enterprises. Therefore, it is vital to keep the data centres operating with no downtime at all. To achieve this, it is required to provide sufficient power to operate the data centre equipment and provide an effective cooling method to maintain its temperature within the safe limits [1-3].

The continuous developments in the processing power of data centres, and particularly the increased heat dissipation requirements from the chips can have a negative impact on the operational performance and the reliability of the IT equipment [4, 5]. The greater requirement for heat dissipation has increased the cooling demand which in turn increases the energy consumption of the cooling system. A number of studies have been undertaken to reduce the power consumption of data centres and improve the operational performance by focussing on four interconnected aspects which are: IT devices and software, power supplies, cooling techniques and general design of data centres [6]. Therefore, thermal management and energy performance enhancements of data centres is an ongoing research problem

to reduce energy consumption in data centre cooling and provide the required thermal environment for efficient and reliable operation.
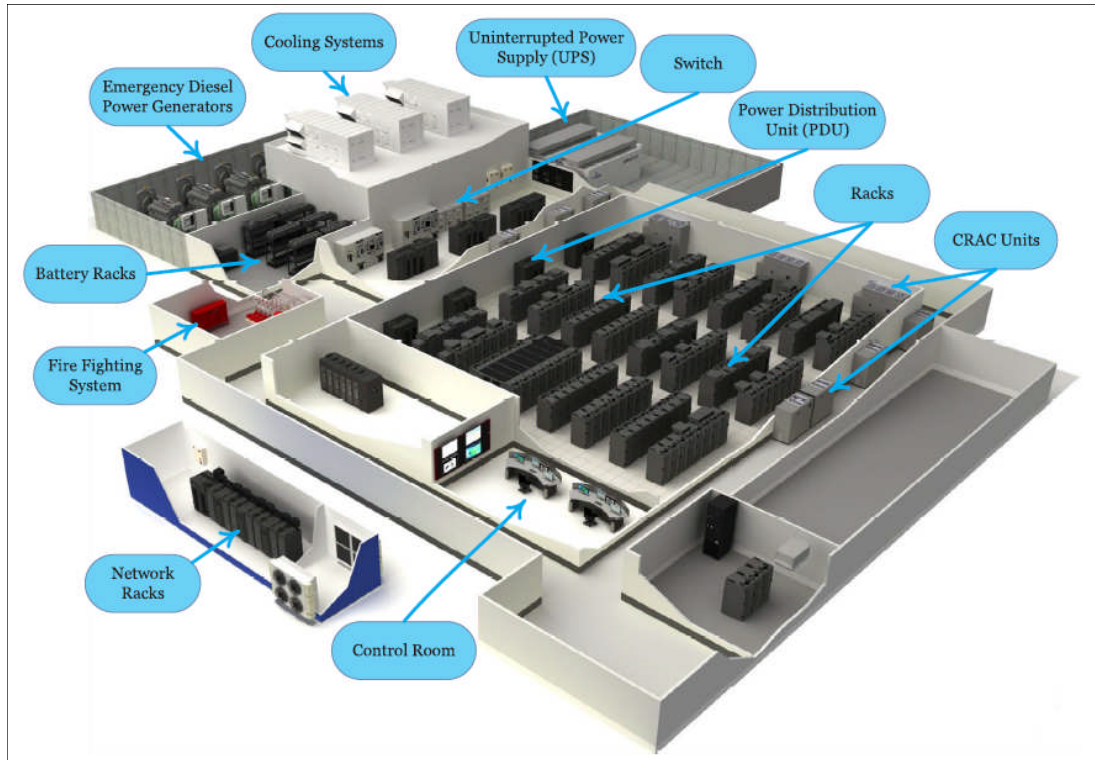
Figure 1.1 Typical data centre infrastructure layout [2].

1.2 Power consumption and energy efficiency of data centres

The revolutionary technologies of data storage and processing applications and the steadily increasing demand from online websites, have been continuously contributing in increasing the power consumption of data centres for more than a decade [7]. For example, the worldwide power consumption of data centres in the year 2005 was twice that of the consumption in the year 2000. By 2010 this increased by a further 56% to constitute 1.3% of the global power consumption [8]. The power consumption of data centres is expected to grow up 15 to 20% every year [9].

In a breakdown of energy consumption in data centres, the energy consumed in the actual IT equipment of racks and its contents is only about half of the power supplied to the whole data centre [8, 10]. However, the remaining

power is considered as the overhead losses which is spent on cooling, electrical power delivery losses, lighting, and other facilities. The electrical delivery losses include the voltage losses due to conversions, uninterrupted power supply units (UPS) losses and power transportation losses within the data centre [11]. Among these overhead losses, the power spent on cooling systems of a typical data centre constitute the largest portion of power consumption which is reported by many publications to be 35% of the total power consumption of a data centre [8, 11, 12]. Figure 1.2 shows the power distribution for the different parts of a typical data centre.
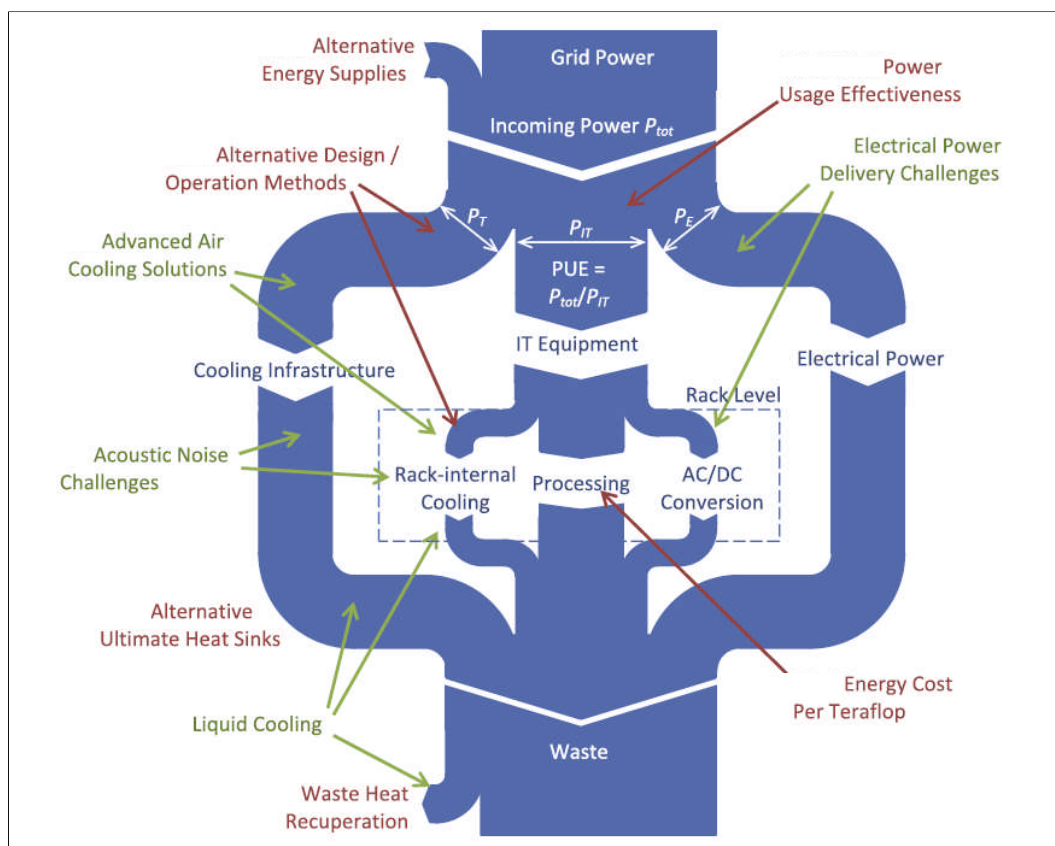
Figure 1.2 Distribution of power into the three main components of a typical data centre (computing, cooling and power losses) [11].

The demand for improved energy efficiency in data centres has raised the necessity to optimise existing technologies and innovate new methods on both the IT side and the supporting infrastructure. A focus of this has been the cooling infrastructure, which represents the largest service requirement.

3

1.3 Multi-level thermal management data centres

The operation of the IT equipment in data centres results in converting almost all the electrical power into heat. This heat is generally transported through multiscale subsystems from the heat generation to the environment as shown in figure 1.3 [1]. It is mainly generated at the chip level of nanometre scale from the transistors and the interconnects, to the server level of the metre scale, to the rack level of metres in scale to the data centre level at hundreds of metres before been transferred to the environment at the facility level [13]. This series of multiscale subsystems of transferring the heat has made the thermal management of data centres challenging for the designers and engineers as the cooling is initiated at the facility level and then conducted back to take the heat away from the chip level.
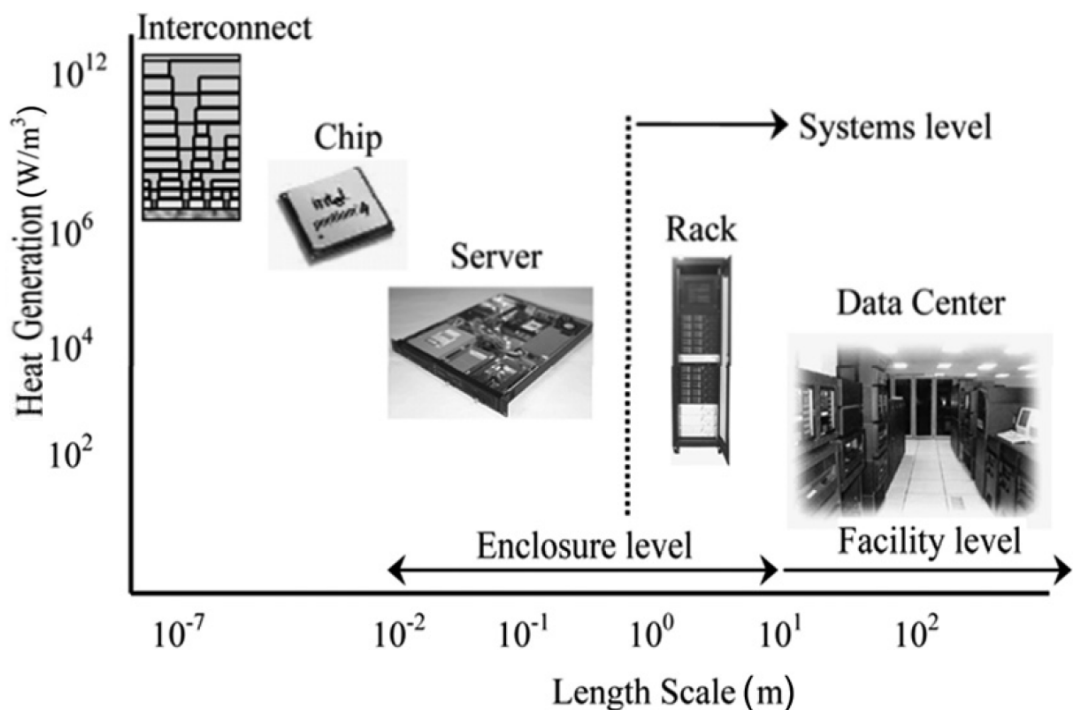
Figure 1.3 Thermal management of data centres through multiscale systems [14].

4

The microsystems packaging have focused mainly at the chip, server and rack levels while the data centres and facility levels have been the focus of heating ventilation and air conditioning (HVAC) engineers [1]. Therefore, holistically studying the cooling system from the chip to the environment is necessary to minimise the power consumption of cooling especially with the high increase in the IT equipment density.

Datacom systems use different approaches to remove the heat and the most common is to use air, this can be done in numerous ways such as underfloor plenums, overhead ducts, containment and side fed [2]. However, by the year 2020, it is expected that the heat generated by the general purpose processors will increase to 190W [15]. Moreover, the generated heat by the high-performance computer processors (HPC) is already as high as it was recorded to be 300W in 2012 [15]. It is clearly shown that the air cooling has reached its limitations of heat transfer as the effective convection coefficient of air ranges between 0.04 to $0.07\frac{W}{cm^2.K}$ with the difference between the inlet and outlet temperatures across the rack of 15℃ [16]. Therefore, it is necessarily to reintroduce liquid cooling as an efficient heat transfer medium to cool the IT as it had been already used in electronics cooling prior to introducing the Complementary Metal Oxide Semiconductor transistor (CMOS) in the 1990s as shown in figure 1.4 [17].

The increasing heat dissipation within data centres and the projected development trend has seen a reconsideration of liquid cooling technologies. Concerns raised with using liquid cooling over traditional air cooling include leakage and complicated infrastructure requirements possibly increasing costs [18]. The principle of the design of liquid cooling systems is to bring the liquid medium closer to the IT enclosure. This can be achieved using different ways such as rear door heat exchanger (RDHx), or bringing the liquid in direct contact with the chip such as the cold plate direct liquid cooling or it can go even beyond that by immersing all the electronics in a dielectric liquid [2].
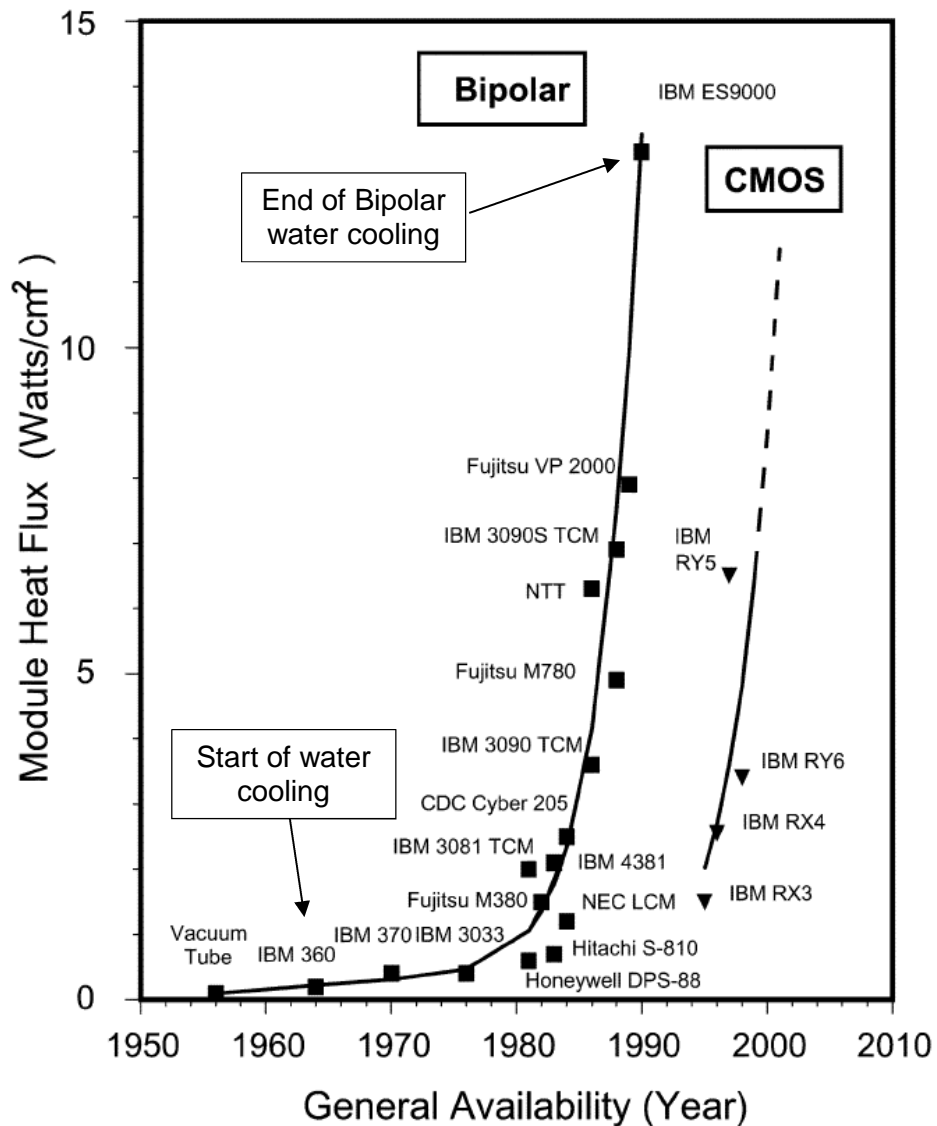
Figure 1.4 History of the module heat flux in the Bipolar and CMOS transistors [17].

The facility environment tiers and the combined boundaries of liquid cooling in data centres define the IT equipment operation limits. ASHRAE [19] described the thermal guidelines of liquid cooling and categorised them into five classes depending on the type of facility cooling equipment. These classes are ranging from W1 to W5 and each class has its own temperature boundaries as shown in table 1.1. Classes W1 and W2, use chiller and cooling tower (depending on the geographical area whether it uses economiser for energy improvements or not). Class W3, typically uses the cooling tower for most of the time and uses chiller under certain conditions. Class W4, uses

only a dry cooler or cooling tower to improve the energy performance of the data centre. The supply water temperature in the W4 class ranges from 2 to 45℃ (which will be utilised in the current study). Finally, W5, completely based on the heat reuse where the data centre dissipates its heat to cool buildings or to be used in any other facilities. The operation environmental class should be chosen based on non-failure conditions with full operation of the IT under the specific environmental conditions.

Table 1.1 ASHRAE liquid cooled thermal guidelines [20].

Class	Typical Infrastructure Design		Facility Supply Water Temperature, °C (°F)
	Main Cooling Equipment	**Supplemental Cooling Equipment**	
W1	Chiller/cooling tower	Water-side economizer	2 to 17 (35.6 to 62.6)
W2			2 to 27 (35.6 to 80.6)
W3	Cooling tower	Chiller	2 to 32 (35.6 to 89.6)
W4	Water-side economizer (with dry-cooler or cooling tower)	N/A	2 to 45 (35.6 to 113)
W5	Building heating system	Cooling tower	>45 (>113)

1.4 Aim and objectives of the current study

To improve the energy consumption and optimise the computational performance in Direct Contact Liquid Cooled (DCLC) data centres, it is required to holistically integrate the workload of the IT servers with a systematic thermal management from the chip through the multilevel systems to the environment by the external heat rejection system. For that, the main aim of this research is to gain an in-depth understanding of thermal and energy management of a DCLC data centre. In addition, explore and investigate the novel design of a chiller-less spray evaporative assisted cooled DCLC data centre. Figure 1.5 shows the holistic management of the integration of four different factors on a DCLC data centre. The aims were pursued via a three level approach: server level, rack level and facility level as shown in figure 1.6
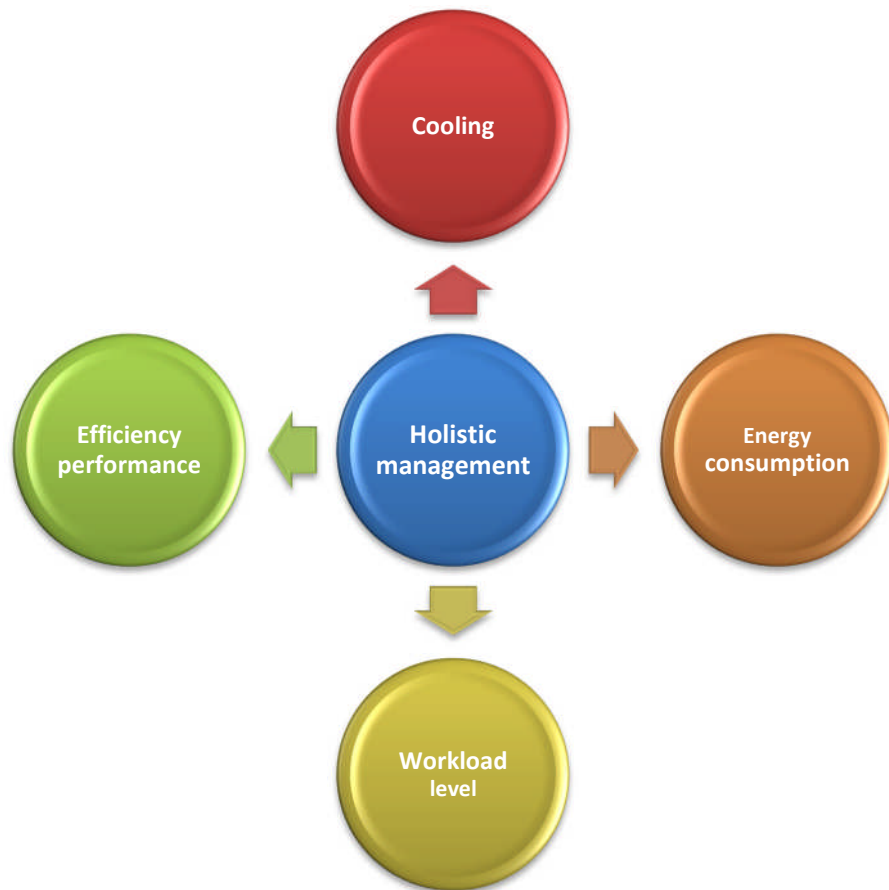
Figure 1.5 The holistic management strategy followed in this study
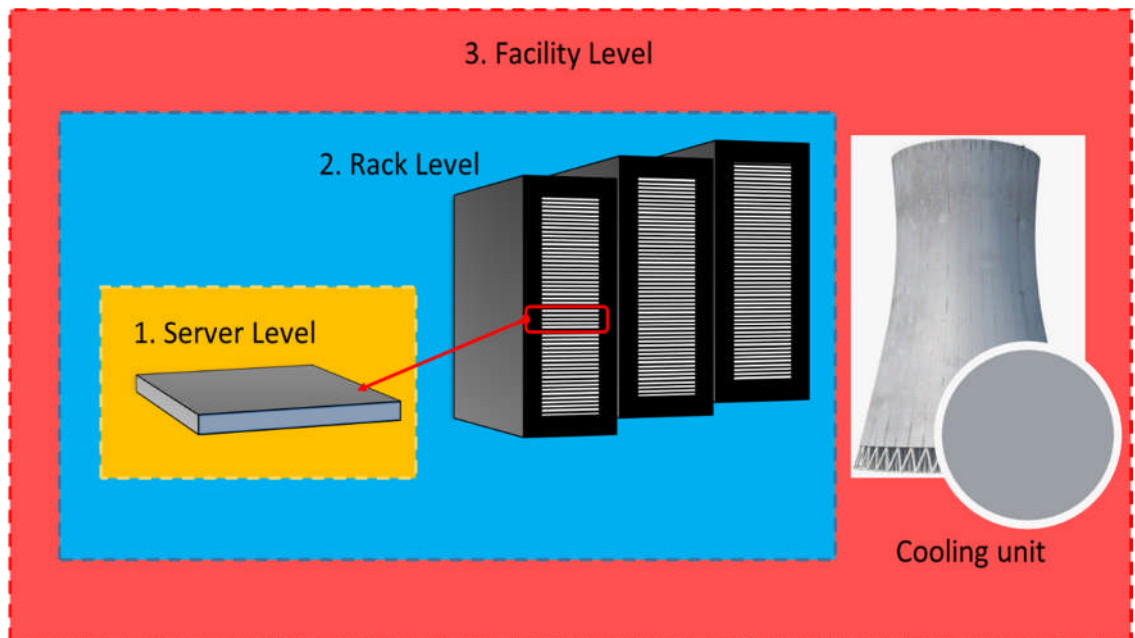
Figure 1.6 The three levels approach to study the holistic management

The following work plan has been adopted in this thesis to achieve the aims of the present study:

1. Design, build and commission a full-scale system of direct contact liquid cooled servers as shown in figure 1.7. This includes 30 Sun Fire V20z servers arranged in a rack level. The servers are hybrid (water/air) cooled using DCLC cold pates and Rear Door Heat Exchanger (RDHx). The design also consists of Air Handling Unit (AHU) that utilise spray evaporation to increase the performance to reject the generated heat by the rack to the environment.

2. Based on the design, consider characterisation of each component in the loop from the server level to the facility level regarding power consumption and thermal performance. This also includes characterising the operational performance of the design for different ambient conditions and workload levels regarding its impact on the servers' component temperatures, power consumption and IT performance.

3. Investigate the impact of rack flow rate, pumping configuration (central and distributed pumping) and flow distribution to the servers on the IT performance, power consumption and servers' component temperatures for different IT workloads.

4. Explore the effect of rack inlet temperature in terms of electronic component temperatures, IT performance and power consumption for a wide range of IT workloads.

5. Finally, carry out a comprehensive investigation on droplet evaporation over heated surfaces of different wettability. This study can help to understand the evaporation mechanism on surfaces that have the potential to be applied in cold plate heads at the server level or in the AHU heat exchangers that use evaporation to increase the heat transfer capacity.
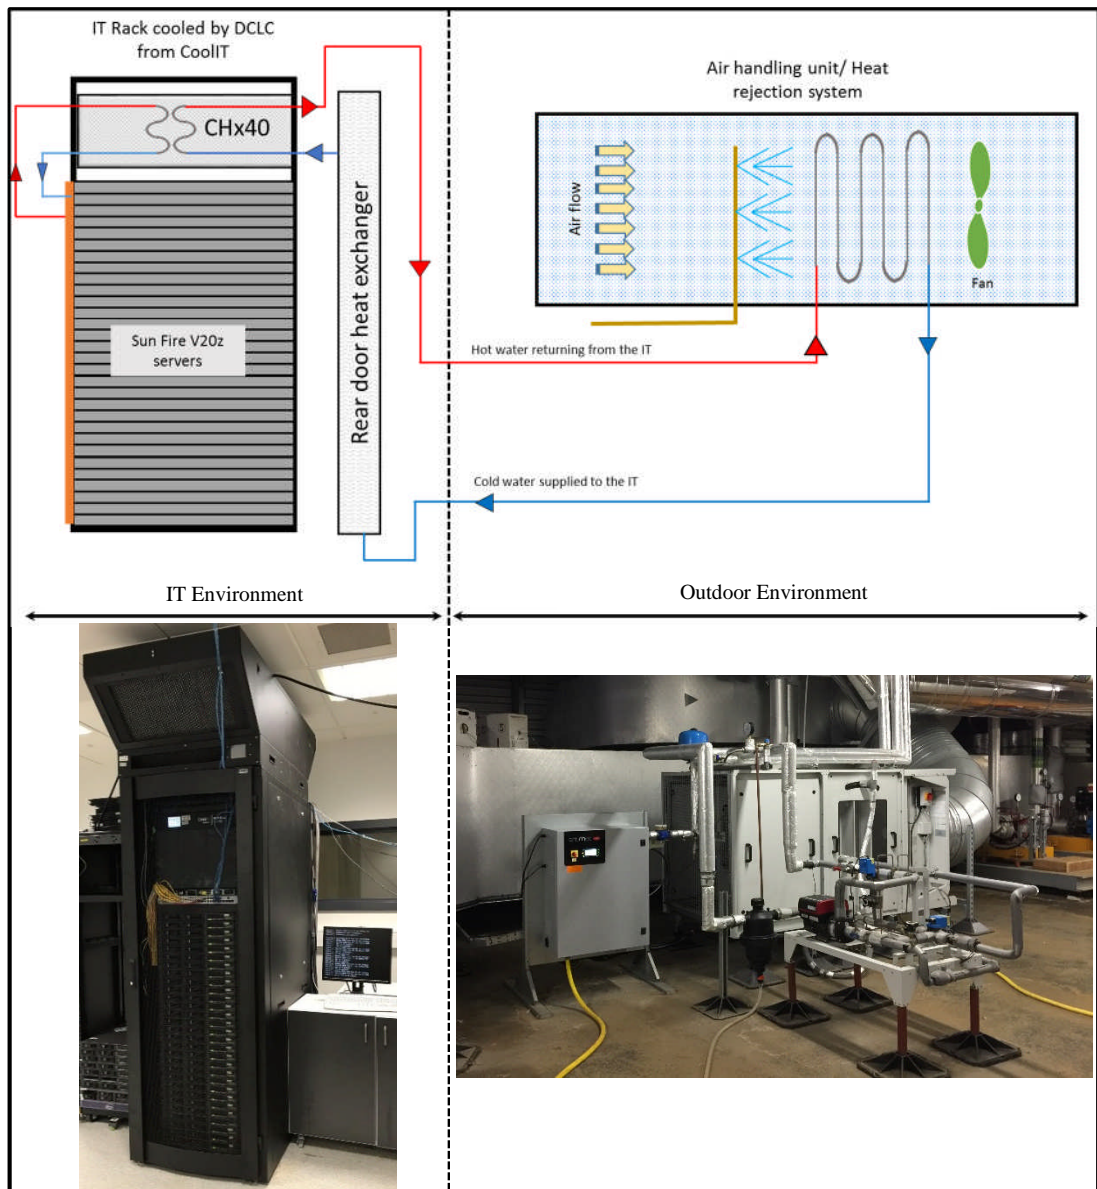
Figure 1.7 Diagram of the DCLC system design configuration

1.5 Thesis outline

To achieve the aims and objectives of this work, the thesis is divided into 9 chapters. These chapters are briefly outlined as follows:

Chapter 1 gives a general overview of data centres and concepts, thermal management challenges in data centres, the aims and objectives of the study and thesis outline structure.

Chapter 2 contains the literature review, which starts by giving a general introduction to data centres and the efficiency metrics, together with an overview of the available cooling technologies. The related studies regarding liquid cooling are reviewed and divided into three categories: server level, rack level and facility level. Finally, the liquid cooling designs overview is given before the summarising section of the chapter.

Chapter 3 describes the design and construction process of the chiller-less liquid cooled data centre starting from the chip to the environment, including extensive commissioning of the unit.

Chapter 4 gives a description of the experimental setup and measurements methodology.

Chapter 5 investigates the operation performance of the design, as well as presenting the characterisation of the system both holistically and for each component.

Chapter 6 investigates the flow distribution inside the rack, pumping configuration and the effect of flow rates. In addition, the impact of these factors on the power usage effectiveness, IT efficiency and thermal performance of the rack are assessed.

Chapter 7 investigates the effect of inlet coolant temperature to the rack on the holistic performance.

Chapter 8 stands alone in investigating the droplet evaporation on different surfaces of heat exchangers. This chapter presents an introduction, literature review, study methodology and the results.

Chapter 9 concludes the finding of the thesis and provides recommendations for future studies.

Chapter 2 Literature review

2.1 Introduction

Thermal management and energy efficiency of data centres have been discussed widely in the literature. However, information relating to direct liquid cooled data centres, particularly the holistic management, still remains sparse. In light of this, the following literature review covers the related topics which are concerned with the operational performance of direct liquid cooled data centres. The chapter will start by reviewing data centre performance metrics, give a general overview for data centre cooling techniques (with attention on direct liquid cooling) and finally an insight into data centre cooling designs will be reviewed.

The published literature that relates to direct liquid cooling will be reviewed in three broad categories: "server level", which focuses mostly on the part of data centre that services the IT load, "rack level" which is the configuration of multiple servers and "facility level" which includes the other facilities of the data centre and more specifically the external heat rejection system (e.g dry cooler or air handling unit).

2.2 Data centre performance metrics

The performance metrics are vital to operating the data centres with low cost of ownership and high operational efficiency and reliability for the purpose of producing the highest computational output as a function of the energy input. In response to these aims, the industry and academia have worked to establish many different performance metrics. Among these metrics, Jamalzadeh and Behravan [21] presented a review of more than 30 metrics proposed by different organisations to evaluate different aspects of data centre energy efficiency. However, organisations have converged on some for measurement and comparison purposes. In this section, the most widely used key performance metrics are discussed. These include thermal management

metrics, energy effectiveness metrics, efficiency metrics, comprehensive metrics, and carbon emissions and water usage metrics.

2.2.1 Thermal effectiveness metrics

To assess the efficiency of individual components in data centres in terms of their ability to deliver the required service, it is necessary to use thermal management metrics. Particularly, the cooling system needs to be thermally assessed as it is a vital part of data centres where an efficient cooling system means less total power consumption and more reliable computational operation.

To evaluate the performance of the cooling system in terms of the efficiency of power input, the Coefficient Of Performance (COP) is a commonly used metric [22]. It is defined as the ratio of the cooling capacity (Q_c) to the power input (W) of the cooling system:

$$COP = \frac{Q_c}{W}$$
2.1

Sharma et al (2002) [23] proposed two dimensionless metrics to evaluate the performance of the cooling system based on the fluid flow and energy balance from rack level to data centre level. The first metric is called Supply Heat Index (SHI) which is defined as the indication of the recirculation effect due to mixing between the cold air and hot air inside the IT room. The SHI defines the convection heat transfer inside the data centre building. The ideal value of the SHI is zero, with no air mixing between the hot and cold aisles. However, higher SHI means there is an over mixing or higher supply air than the required. The second metric is called the Return Heat Index (RHI) which is defined as the ratio of the total heat extracted by the computer air conditioning unit (CRAC) to the enthalpy change at the rack exhaust. The change in the enthalpy at the exhaust of the rack is due to infiltration (air mixing) which results in the algebraic summation of the SHI and RHI metrics being equal to unity.

Herrlin (2005) [24], developed a metric called rack cooling index (RCI) which gives an indication of how effective is the data centre cooling system to

maintain the inlet temperatures to the racks within the recommended values. The RCI factor is defined as follow:

$$RCI = \left[1 - \frac{Total\ Over\ Temp}{Max\ Allowable\ Over\ Temp}\right] 100\%$$ 2.2

Where $Total\ Over\ Temp$ is the number of servers in a data centre which have temperatures higher than the recommended and $Max\ Allowable\ Over\ Temp$ is the number of servers which have temperatures higher than the maximum allowable.

Another metric was also developed by Herrlin in 2007 [25], and is called the return temperature index (RTI). The RTI is a measure of the infiltration that may occur in the IT room between the cold air and hot air through the cold and hot aisles. The RTI is defined as follows:

$$RTI = \left[\frac{T_{Return} - T_{Supply}}{\Delta T_{Equip}}\right] 100\%$$ 2.3

The RTI is an important metric, since it assesses how effective the air distribution is inside the data centre. The ideal value of the RTI is 100%, where lower than this means there is air bypassing the servers and waste energy while higher values means that air is mixing and a deficiency in air separation occurs.

Vangilder et al (2007) [26] developed a new metric to calculate how effective the racks are arranged with respect to the CRAC/CRAH units in the room. The metric is called capture index (CI) and has two versions. One version for the cold aisle and defined as the fraction of air which is captured by the racks to that supplied by the cold air sources. The second version is for the hot aisle, which is defined as the fraction of air received by the local extracts to that exhausted from the racks.

Exergy analysis is one of the powerful metrics that can be used to evaluate the performance of data centres by calculating the ability to produce useful work. The exergy metric is based on the second law of thermodynamics, unlike other metrics which are based on the first law. The first law metric deals with the efficiency of the used energy. However, the second law of

thermodynamics metric considers how the energy is transferred. Lettieri et al [27] compared between using the exergy metric with the first law metrics. They found that the benefits of using the second law metric overweighed the first law metrics as it takes the environmental impacts into consideration. Brady [28] highlighted that the exergy metric is a very useful way to evaluate the performance of data centres, particularly when it comes to heat reuse studies as it provides a case for heat reuse. However, there is complexity associated with the process of measurements and calculations compared with the first law metrics, making it challenging to apply.

2.2.2 Energy effectiveness metrics

The most commonly used metric to describe the energy effectiveness and usage in data centres is the Power Usage Effectiveness (PUE). This metric was first proposed by Malone and Belady in 2006 [29] and further popularised in 2007 by the Green Grid group [30]. PUE is defined as the ratio between the total power required to operate the data centre and the power required to operate the IT equipment. Another form of PUE was also developed by the Green Grid organisation and is called Data Centre Infrastructure Efficiency ($DCiE$) which is the reciprocal of the PUE.

$$PUE = \frac{1}{DCiE} = \frac{Total\ power\ supplied\ to\ the\ data\ centre}{Power\ supplied\ to\ the\ IT\ equipment} \qquad 2.4$$

The PUE shows the proportion of energy which is actually used to operate the IT to the total power drawn by the data centre facility. Thus, a data centre of lower PUE value has better effectiveness and the ideal effectiveness of being 1. A data centre with a PUE of 1 indicates the ideal case that all the power supplied to the data centre is consumed by the IT equipment [31].

The concept and mathematical simplicity of PUE have made it widely applicable to assess the effectiveness associated with the overheads for data centres. However, a number of studies have raised a concern of using it by its own to represent the performance of data centres. Brady et al (2013) [32] critically assessed calculating the PUE in data centres by conducting a case study to analyse the effectiveness of the social network company Facebook data centre at Prineville. The study highlighted the level of detailed

engineering data which is required to calculate proper PUE values of a data centre. Their investigations showed that the PUE is crucial to encourage energy usage reduction. However, it cannot be considered a clear metric for energy efficiency. Garimella et al. (2013) [11] reported that the PUE is restricted in covering the cooling components inside the servers (e.g servers small fans) as it considers all the power consumed to the IT at the rack input as the IT component power consumption. They also argued that the PUE does not include the operational expenditure that supports the data centre's operations such as water treatment used for cooling. Horner and Azevedo (2016) [33] highlighted three main issues associated PUE which make it an incomplete metric even though its usefulness in measuring the data centre facility overhead. The first issue is associated with energy efficient viewpoint, where it is unable to capture the hardware efficiency of the IT. The second issue is that the PUE cannot evaluate the productivity of the data centre which is associated with the business point of view. The third shortcoming of the PUE comes from its limitations in measuring the carbon emission with respect to the energy performance from the environmental considerations. Despite these shortcomings, the PUE metric has provided the sector with improvements and recently the PUE metric has become an ISO standard with clear guidelines on how the metric should be obtained [34].

Another metric found in the literature is based on the ratio of power consumption to floor surface area which is called Power Density. This metric is used to assess the space usage of a data centre mostly for the economic considerations [28]. However, this metric has limitations on its applications and cannot be used to compare different data centres as it is based on the surface area that is measured so it is different locally inside one data centre as depicted by Sun and Lee (2006) [35].

Sun and Lee (2006) [35] developed a metric based on a case study on data centre energy performance called $Relative\ Energy\ Effectiveness$. This metric has two terms, the first one is for the IT which is similar to the PUE, and the other one is for the heat ventilation and air conditioning system (HVAC) as follows:

$$IT_{effectivness} = \frac{IT\ power\ consumption}{Total\ power\ consumption} \qquad\qquad 2.5$$

$$HVAC_{effectivness} = \frac{HVAC\ power\ consumption}{IT\ power\ consumption} \qquad\qquad 2.6$$

Tschudi et al. (2010) [36] proposed a modified metric of PUE to accommodate the heat reuse energy. The metric is called energy reuse effectiveness (ERE) which is similar to PUE apart from the power consumption that is subtracted from the numerator as follows:

$$ERE = \frac{P_{Total} - P_{reuse}}{P_{IT}} \qquad\qquad 2.7$$

On the modifications of the $DCiE$ metric, Shash et al. (2011) [37] modified the definition of the effectiveness by considering the power consumed inside the server by the leakage, fans, and idle power part of the total power in the $DCiE$ equation. The idea of the metric is to include the computation power consumption in the metric as it is considered as the useful output of the data centre. This metric is called the grand coefficient of performance(COP_G):

$$COP_G = \frac{P_{IT} - P_{IT,leakage} - P_{IT,fans} - P_{IT,idle}}{P_{Total} + P_{IT,leakage} + P_{IT,fans} + P_{IT,idle}} \qquad\qquad 2.8$$

The boundary between the servers' infrastructure and facility infrastructures has been a subject of debate in PUE calculations. For example, it is not clear on which side the power consumption of the servers' fans and power supply losses should be accounted for. For this, Patterson et al. [38] proposed a new metric of effectiveness for the IT equipment which is called IT-power usage effectiveness ($ITUE$). The $ITUE$ is described as the total power consumed by the IT equipment to the power consumed by the IT to perform the computations. The proposed $ITUE$ metric is the analogy of the PUE but for the IT equipment and not the data centre. In another words, the PUE focuses on the infrastructure of the data centre while the $ITUE$ focuses on the infrastructure of the IT equipment. Figure 2.1 shows the schematic of the data centre and IT equipment. From this figure the PUE and $ITUE$ can be calculated from equations 2.9 and 2.10, respectively, as follows:

$$PUE = \frac{a+b}{d} \qquad\qquad 2.9$$

$$ITUE = \frac{g}{i} \qquad\qquad 2.10$$

The combination of PUE and $ITUE$ give a new metric called the total usage effectiveness (TUE):

$$TUE = PUE \times ITUE = \frac{a+b}{i} \qquad\qquad 2.11$$

The same authors also carried out case study calculation to assess the $ITUE$ and TUE metrics on the ORNL's Jaguar data centre. The study stands for comparisons between two types of servers in a data centre one with higher efficiency and the other one with lower efficiency. The results showed that the low-efficiency servers have lower PUE compared to the high-efficiency servers. However, the $ITUE$ and TUE are shown to be lower for the high-efficiency server. These results showed that the low-efficiency data centre has a higher power consumption even though they have the lower PUE. They concluded that the PUE may give a false indication if it is used to compare the effectiveness of two data centres. This means that a data centre of lower PUE value does not necessarily mean that it is more energy efficient. However, the $ITUE$ and TUE give more reliable results for the comparison.
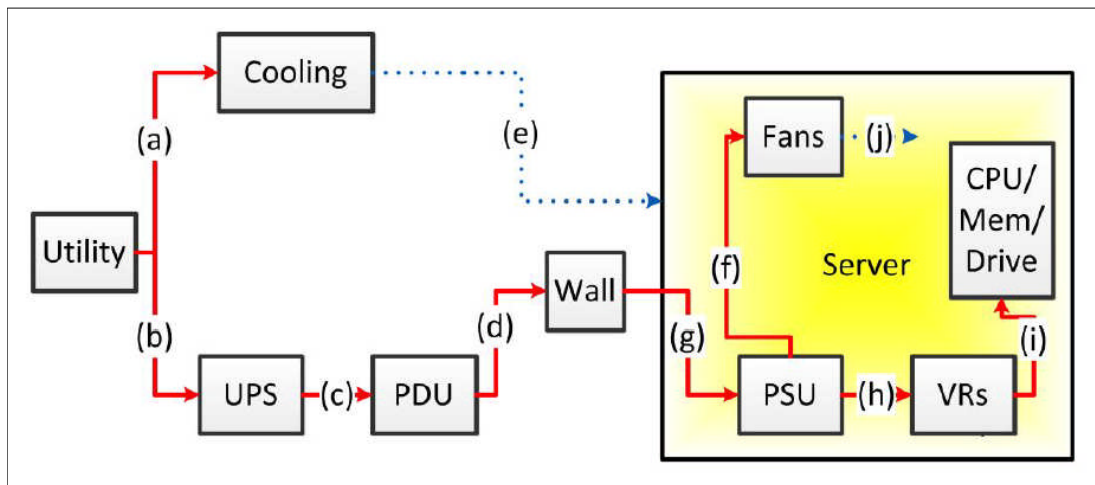
Figure 2.1 Schematic explanation of data centre and IT equipment to show the points where the PUE, $ITUE$ and TUE are taken [38].

2.2.3 IT efficiency metrics

The efficiency of the IT equipment in terms of power consumption and data processing is vital to be captured to draw a complete picture of data centre efficiency performance. To achieve this, it is necessary to look closer inside the servers and analyse the operation under different circumstances.

The most prominent metric that is used to assess the compute efficiency of the server is the $FLOPS/W$. this metric is the most ranked of energy efficient HPC in the Green500 list [39]. It is also called a $Performance/W$ metric which is the ratio between the server floating point operation per second and the power consumption of the server as shown in equation 2.12. The higher efficiency of the servers comes with the higher $FLOPS/W$ the system can deliver.

$$\frac{FLOPS}{W} = \frac{Average\ achieved\ performance}{Average\ IT\ power\ consumption} = \frac{Floating\ Point\ operation\ per\ second}{W} \qquad 2.12$$

Minartz et al. (2010) [40] defined a new metric to evaluate the energy efficiency of HPC servers termed Energy-to-Solution (ETS). The ETS gives an indication of the energy consumption of HPC servers to carry out a specific computation using a specific software application.

Newcombe et al. (2012) [41] introduced a metric to target the wasted energy during the idling operation of the servers where there is no useful work produced. This metric is called Fixed to Variable Energy Ratio ($FVER$) and it is defined as follows:

$$FVER = 1 + \frac{Fixed\ Energy}{Variable\ Energy} \qquad 2.13$$

The $Fixed\ Energy$ represents the power consumption of the servers when it is idle and the $Variable\ Energy$ represents the power consumption when the servers perform useful computational operations. The range of $FVER$ is between 1 and 10. The ideal $FVER$ with being 1 which means that the servers consumed no power when they are not producing any useful work which is practically not possible.

Subramaniam and Feng (2012) [42] proposed a metric called The Green Index (TGI) to capture the energy efficiency of supercomputers. The metric uses different benchmark tests to combine different performance outputs and provides a unique number to represent the energy efficiency. The authors presented a methodology of calculating the TGI which can be achieved in three steps. The steps start with calculating the performance to power ratio then obtaining the relative energy efficiency (REE) from the benchmark and calculating the weighing factor (W) which is based on all the benchmarks. Finally, the TGI is calculated by summing all the $REEs$ and their corresponding W as follows:

$$TGI = \sum_i W_i \times REE_i \qquad\qquad 2.14$$

Beitelmal and Fabris (2014) [43] developed two metrics for efficiency measurements at the server level and data centre level based on the thermodynamic approach. The metric for the server level is called server energy performance productivity ($SEPM$) and the one for the data centre level is called data centre energy performance productivity ($DCEPM$). The metrics are defined as the ratio of the dynamic power consumed due to the utilisation of the server (or servers for the $DCEPM$) and the total system power consumption.

2.2.4 Comprehensive efficiency and effectiveness metrics

Assessment of the data centre infrastructure effectiveness and IT performance efficiency are key steps towards the efficient operation of data centres. However, it is important to integrate between these two metrics to obtain a more generic relationship that can link between the infrastructure effectiveness of a data centre and its output computational performance. The gap between the data centre overhead effectiveness (PUE) and the servers' performance ($FLOPS/W$) make it possible to improve one of them, but the other will deteriorate. For example, increasing the temperature inlet to the rack will have a positive effect on the PUE of a data centre as the overhead decreases which reduces the power consumption required for cooling. However, this may have a negative impact on the internal components of the

IT equipment where the componenet temperatures increase is associated with an increase in the current leakage and in turn increases the power consumption of the servers and decreases the performance.

For this, Wilde et al. (2014) [44] proposed a new metric called data centre workload power efficiency ($DWPE$) to bridge the gap between the data centre infrastructure effectiveness and efficiency. They divided the data centre into four pillars as shown in figure 2.2. The figure shows the clear gap between the PUE and the $Performance/W$ metrics. They introduced two new metrics to drive the $DWPE$ from them. The first metric is called system PUE ($sPUE$) which is similar to the normal PUE. The second metric is called workload power efficiency (WPE) which is similar to the FLOPS/W metric. It should be mentioned that the WPE in their study include the power consumption of the complete system. Hence, they came up with the new metric which can be used to calculate the efficiency performance of a data centre for a specific computational load which is defined as the ratio between the WPE and the $sPUE$, as follows:

$$DWPE = \frac{WPE}{sPUE} \qquad\qquad 2.15$$
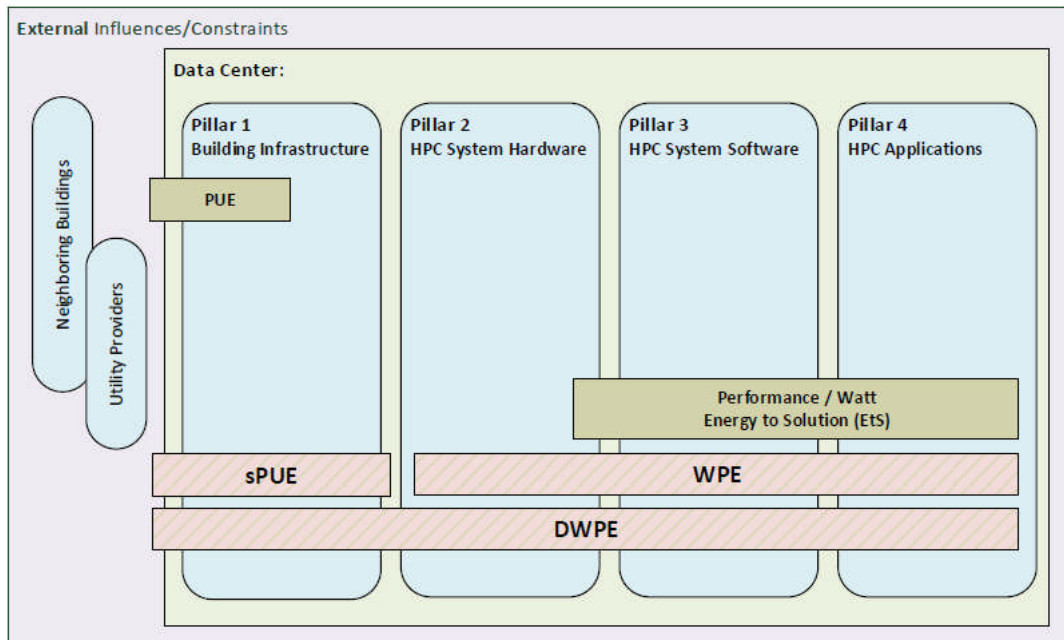
Figure 2.2 Schematic explaining the proposed metrics by Wilde et al. and the covered areas in the data centre [44].

2.2.5 Environmental concerns and water usage effectiveness

The environmental impact of data centres in terms of carbon emissions and water consumption is necessary to be assessed as important as the other metrics of energy and performance. A study by Beitelmal and Fabris (2014) [43], stated that 1kWh of power consumed of electricity generates about 0.87Kg of CO_2. They also stated that data centres' carbon emissions constitute about 2% of the global carbon emissions. In line with the carbon emissions, data centres consume a large amount of water. For power, water consumption is not specific to data centres, however in warm and dry climates, water is commonly used in evaporative and adiabatic cooling and so water usage in these situations where water scarcity may be a problem itself. The water consumption in data centres is highly dependent on the means of power generation and cooling technology.

These indicators drive effectiveness metrics to measure the usage of carbon and water usage in data centres. Water usage energy (ω) is also another form of effectiveness metrics to calculate the sustainability of data centres which was introduced by Sharma et al. (2009) [45]. They developed this metric to capture the water energy which is defined as the ratio of the energy embedded in the water (footprint) to the power consumed by the IT (E_{IT}). The power embedded in water is defined as the energy required to distribute and treat the water to the demand location which is divided into two sections: direct usage which is the water used for cooling of data centres (E_d) and indirect usage which is used in the power generation plants (E_n).

$$\omega = \frac{[E_d + E_n]}{E_{IT}} \times 10^3 \qquad\qquad 2.16$$

Azevedo et al. (2010) [46] proposed a metric called the carbon usage effectiveness (CUE) which uses the IT energy consumption as a reference in the same way of calculating the PUE. The CUE is defined as the ratio of the carbon emission in a data centre to the IT power consumption as follows:

$$CUE = \frac{Total\ CO_2\ emissions\ caused\ by\ the\ total\ data\ centre\ energy}{IT\ equipment\ energy} \qquad\qquad 2.17$$

2.3 Thermal management of data centres

The heat generated by servers is due to the operation of the electronics including central processing units (CPUs), memory modules (RAMs), power supply units, voltage regulators, and chipsets. In terms of thermal management, the CPUs are generating the largest portion of the total heat generated by the servers in relatively small areas compared with the remaining components of the server [16]. Consequently, several studies have been found in the literature that focuses on the CPUs from the geometry and heat loads perspectives.

The cooling solutions of servers can be categorised into three main groups based on the cooling mediums: air cooling, direct liquid cooling and total liquid cooling. Figure 2.3 summarises the cooling techniques which were found in the literature.
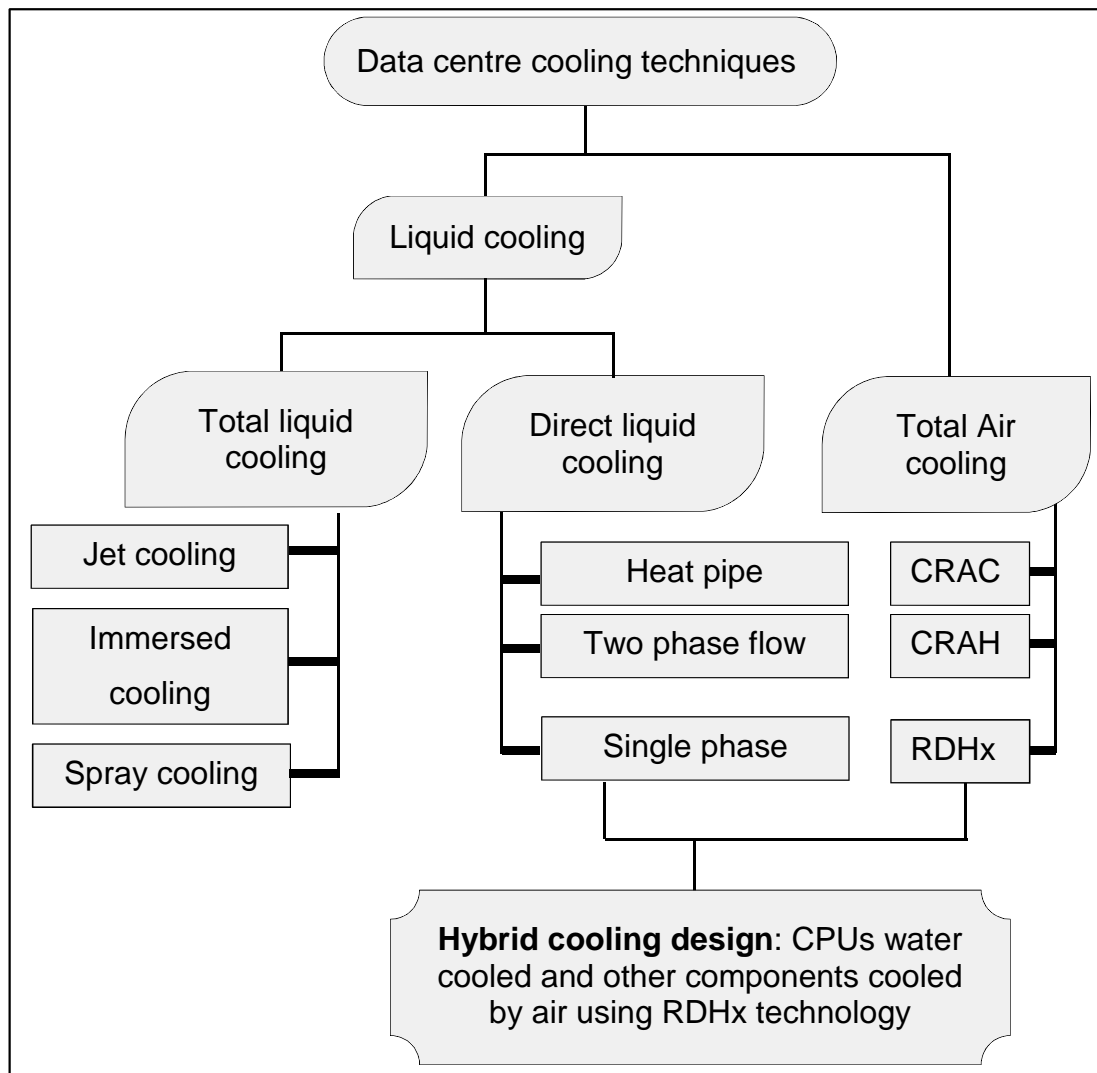
Figure 2.3 Data centres cooling solution methods

2.3.1 Air cooling

Air cooling techniques are the most commonly used and well-established methods in the data centre industry. Typically, the racks housing the servers are arranged in rows which are divided by aisles. These aisles are either for air intake which are called cold aisles or for the air exhaust which are called hot aisles [2, 47]. Figure 2.4 shows the air cooled rack configurations placed over a raised floor to allow for the cold air passages. Chilled air passes in the air passages under the raised floor to pass through perforated tiles in the cold aisle to cool the servers. This air passes through the servers to cool the

electronic components with the help of small fans inside the servers then exhausted from the back of the racks to the hot aisle side [2, 16, 47].

The exhausted hot air from the hot aisle is cooled by Computer Room Air Handling Unit (CRAH) or by Computer Room Air Conditioning (CRAC) [48]. The CRAH utilises water/air heat exchanger to cool the air while CRAC utilises the refrigerant evaporators to cool the air [48].
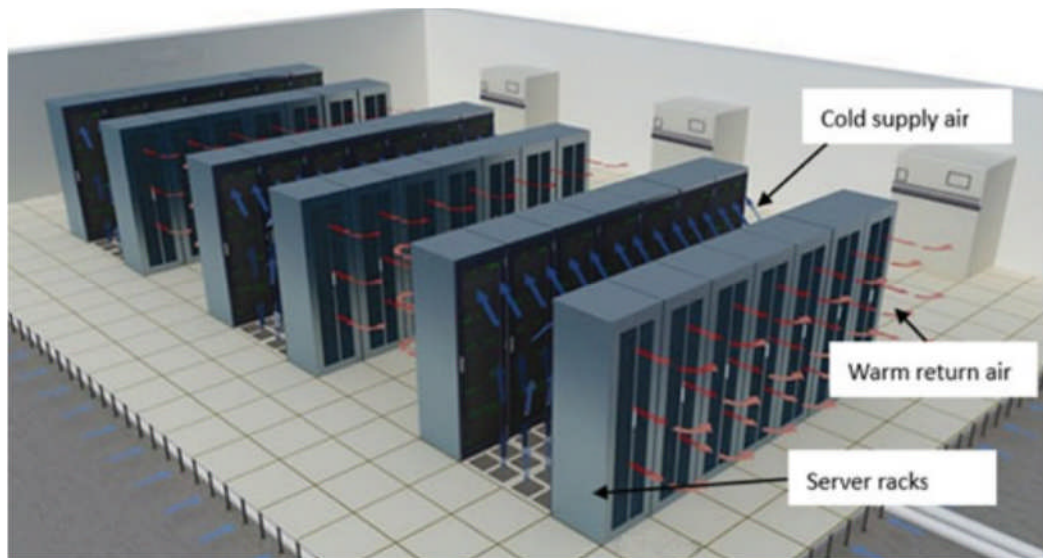
Figure 2.4 Air cooled data centre arrangements [47].

The air cooled data centre facility consists mainly of three parts: the external heat rejection system which is the refrigeration chiller plant, building water pumps, and rooms CRAH or CRAC units. For this, the cooling system consists of about half of the data centre power consumption including 33% for the compressor of the external heat rejection system and about 9% for the room air conditioning unit which is responsible for the air movement inside the data centre as shown in figure 2.5 [47, 49].

The operational nature of data centres is usually dynamic and unpredictable which makes the thermal management a challenging issue for the designers. As a result of that, the air cooled data centre are thermally managed and designed largely by acquired experience and intuition. However, providing proper flow rate at proper temperature inlet for different operating conditions of data centre are the two aims of successful thermal management of air cooled data centres [2]. Several negative effects on the operation of the data

centre result from poor thermal management. For example, low reliability and air maldistribution, increases the shutdown which are all results of the early failure of the servers and increases in the operating cost [50].

A number of shortcomings associated with air cooling in data centres have been found in the literature [2, 13, 47]. These drawbacks are summarised in the following points:

Chiller	33%
Humidifier	3%
CRAC	9%
IT equipment	30%
PDU	5%
UPS	18%
Switchgear	1%
Lighting	1%

Indoor Data center heat
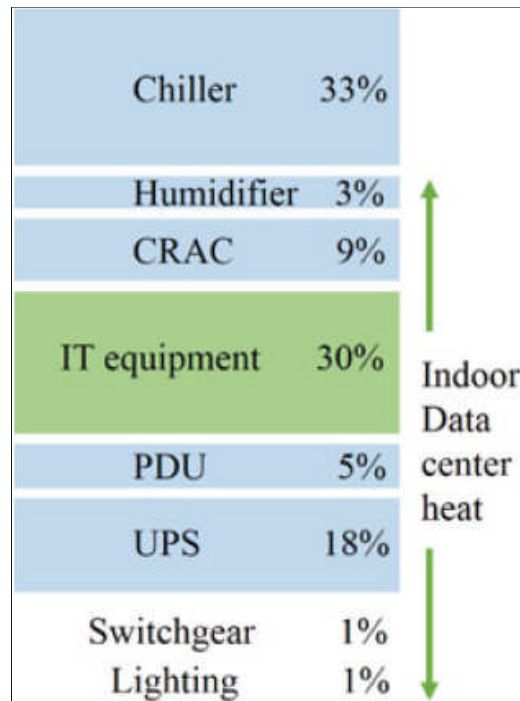
Figure 2.5 Energy consumption by the main components of a typical air cooled data centre [49].

1. Low heat transfer convection coefficient due to the low thermal properties of air requires low supply temperature and high air flow rate which increases the operating costs. The low thermal properties of air are coupled with a high temperature difference of 20 to 30℃ within the servers [2, 16, 48, 51, 52]. For this, a numerical study by Xu et al [53] found that the maximum heat transfer of air cooled heat sink is 340W for a heat sink area of 100cm^2 and CPU die temperature of 90℃. This coupled with the highest heat flux in the heat sink obtained experimentally to be 4W/cm^2 by Katoh et al [54].

2. Due to the high rates of heat generation of some components of the servers, temperature hot spots can be generated. The hot spots can also be generated outside the rack due to the poor air distribution. This necessitates the provision of chilled air with unnecessarily low temperatures to cool the hot spots in the data centre which results in high energy cost required for cooling as well as servers room to be overcooled. This can also result in uncomfortable temperature to the operator [2, 47].

3. The compressor in the refrigerant system needs to be in operation all the times of the year even in the cold environment of the cold regions [2, 47]. This is because of the ambient temperature needs to be sufficiently lower than the required inlet temperature to the servers room to be able to rely on the free cooling [55].

4. High power consumption due to the fans' operations, pumps, and compressors [2, 9, 47].

5. Difficult to harvest and reuse the heat to be utilised for different useful applications as the heat is indirectly collected from the heat generation sources in the servers by air [47].

6. One of the main contributors to air cooled data centre inefficiency is the air mixing. The cold and hot aisles arrangements help to reduce the air mixing but it does not prevent it completely due to cold air bypass and hot air recirculation [23, 56, 57].

The drawbacks of air cooled data centres can be reduced by improving the design of the heat sink over the CPUs. For example, increasing the fin area and heat sink base which both have a positive effect on the heat transfer although it conflicts with the spatial requirements [58]. Using raised floor with ceiling return ducts was investigated by Srinarayana et al. [59] for the 310 kW data centre and found an improvement of $5°C$ in the data centre temperature over the traditional design. Improving the operational strategies can also have an improvement on the operational performance of data centres, for example,

Breen et al. [60] found that increasing the inlet air temperature by 5℃ can improve the coefficient of performance (COP) by 8%. However, this has a negative impact by increasing the hot spots within the data centre. Moreover, using rear door heat exchanger (RDHx) techniques is a prominent method to improve air cooled data centres which will be explained in detail in section 2.5 of this chapter.

2.3.2 Total liquid cooling

In the total liquid cooling technique, the servers are brought in direct contact with the heat transfer fluid where part or all the electronics are covered. This can be achieved in three ways: (i) immerse the electronics in a dielectric liquid, (ii) spray the electronic components; (iii) jet impingement on the hot spots of the servers [16]. Total liquid cooling benefits of keeping the temperature uniform across the whole servers by eliminating the hot spot points. In addition, it offers a versatile cooling solution where no sealed enclosure and piping are a needed [61].

Although total liquid cooling improves the cooling performance of data centres, a number of issues are found to associate with the operation. The thermal properties of water are notably superior compared with dielectric fluids [16]. Therefore, there is a limited heat transfer enhancement compared to the cooling techniques that use water as the heat transfer medium. Using total liquid cooling is also associated with additional required maintenance to prevent dielectric liquid losses due to evaporation [62, 63]. Another difficulty lies in the maintenance and services particularly when swapping electrical components [61].

2.3.3 Direct liquid cooling

Direct liquid cooling is based on removing the heat from electronics by using liquid cold plates or evaporator to replace the conventional air cooled heat sinks [48]. The traditional direct liquid cooling methods are based on using water blocks and cold plates, however, recent studies have focused on enhancing the heat transfer and pressure drop characteristics of microchannel

heat sinks as they have superior features compared with the traditional cooling methods of direct liquid cooling [2, 16, 64-66].

The typical configuration of direct liquid cooling is to use a Coolant Distributed Unit (CDU) to exchange the heat between an external loop and an internal loop. The external loop is connected to a heat rejection system to transfer the heat to the environment while the internal loop is thermally attached to the heat source inside the electronics as shown in figure 2.6 [16, 67]. The liquid coolant is delivered to the racks then the servers and finally to the chip level depending on the used technique of cooling. In most cases, the coolant is only delivered to the microprocessors due to the simple geometry and high heat generation in these components compared with the remaining servers' components which are usually cooled by air [61]. However, some researchers extended the design to cool the other components of the servers with the indirect liquid cooling techniques as will be explained more in section 2.5 [67].

Using liquid cooling extends the bounds of heat transfer and reduces power consumption of data centres due to the high thermal properties of liquids compared to air [58]. The accommodation of direct liquid cooling in data centres is a reduction in versatility because of the requirement of installing piping networks and sealed enclosures required to deliver the liquid coolant to the chip level inside the servers [62]. For this, the cooling system of the data centre needs to be tailored to the servers' position and cooling facility layout to utilise the direct cooling system. On top of that, using direct liquid cooling restricts the easy access to servers which commonly exist in air cooled data centres [62]. However, quick disconnected fluidic connectors (such as stäubli valves) can solve the problem of hot swap access but there exists increases the potential for risk of leakage over the electronics [58, 62].

The direct liquid cooling can be categorised based on the method of cooling at the chip level. The main methods are single phase cooling, two phase cooling and heat pipes [16, 58]. However, the single phase technique will be accommodated in this study through the CoolIT systems server modules or what is called "Direct Contact Liquid Cooling (DCLC)". Therefore, the single

phase of direct cooling technique will be explained in details in the following sections.
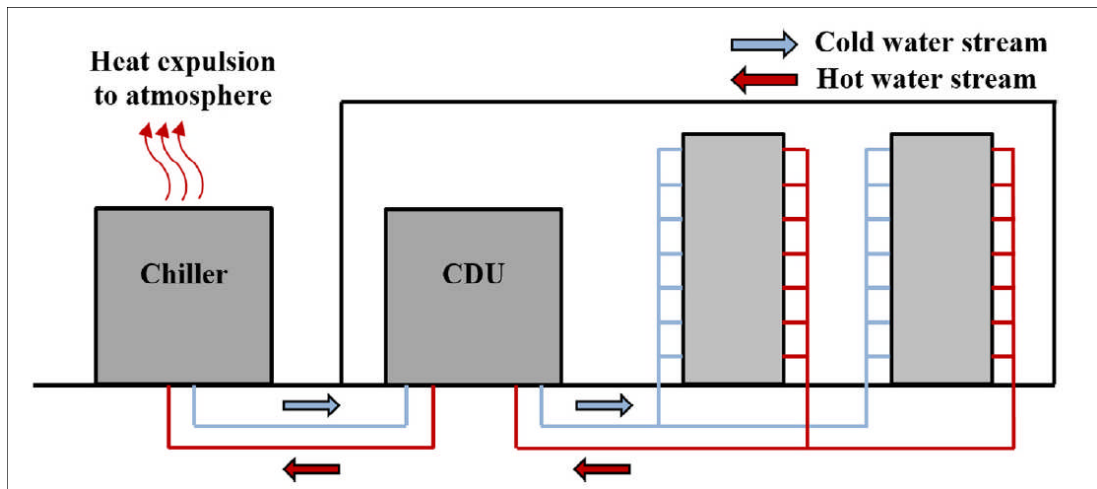
Figure 2.6 Schematic of typical indirect liquid cooling [2].

2.4 DCLC data centres

The generated heat by the IT equipment is removed through the data centre medium. The heat is first generated at a chip level. This heat is then transferred to the liquid cooling loops to be transported to the environment. The process of removing the heat from the IT equipment is generally achieved through three main subsystems, starting with the chip and servers levels, the rack level and the facility level. However, cooling of data centres is started at a facility level and is then conducted back through the three multiscale systems to cool the chip by taking the heat away.

Several thermal management methodologies have been found in the literature in order to save energy and increase the performance of direct liquid cooled data centres. These methodologies are considered in the following sections by categorising them into three groups of studies: chip and server level studies, rack level studies and facility level studies.

2.4.1 Chip and server level studies

Thermal management at the chip level is generally based on two methods: reducing the dynamic power consumption and cooling of the chip. The

dynamic power is based on the dynamic voltage and frequency scaling. This method is utilised in CPUs which results in high power reduction by reducing the operational frequency [68]. Many studies were found in the literature that already applied this technique to control and efficiently maintain the core temperature of the CPUs to obtain significant energy saving [69-71]. The cooling of the chip can be achieved through different methods and techniques as mentioned previously. However, studies on improving the performance of direct liquid cooled electronics at the chip level are basically based on enhancing the microchannel cold plates, flow rates, and inlet temperatures of the coolant.

2.4.1.1 Cold plate heat sink

A high heat transfer coefficient between the cold plate and the CPU surface reduces the temperature difference between the coolant and the cold plate [47]. For this, incorporating microchannels and minichannels in cold plates and on-chip cooling applications have shown a significant improvement in the energy efficiency of these devices [72]. Kandlikar and Hayner [73] provided a detailed description of the different types of cold plate heat sinks and the manufacturing issues. They identified the main factors in the cold plate design as the coolant type, pressure, coolant flow rate, temperature inlet, and type of the material. They also classified the cold plate types by the required power dissipation into four categories (all technical details from [73]):

1. Formed Tube Cold Plate (FTCP).

In this design, the coolant tubes are soldered to the substrate as shown in figure 2.7a. These tubes are either made of copper or aluminium. This design is usually used for low power applications as it provides a simplicity in the manufacturing although it has the lowest thermal performance.

2. Deep Drilled Cold Plate (DDCP)

To reduce the thermal resistance between the coolant and the substrate in the FTCP design to increase the performance, the flow paths of the coolant inside the substrate are created by drilling holes in the substrate itself. The coolant

passage layout is affected by the design of the electronic device. However, it is common to use more than one path of the flow to achieve the required pressure drop, flow distribution and temperature difference as shown in figure 2.7b.

3. Machined Channel Cold Plate (MCCP)

To increase the thermal performance of the cold plate design, the flow paths are machine-cut into the substrate and a lid is placed to form the flow paths for the coolant as shown in figure 2.7c. The flow passage size is varied from 200μm to several millimetres depending on the required heat capacity for different applications.

4. Pocketed Folded Fin Cold Plate (PFCP)

For high heat flux electronic applications, fins are implemented in the cooling passages to increase the heat transfer surface area and increase the local heat transfer coefficient. In this type, the substrate is machined in a way to insert the required type of fins as shown in figure 2.7d. The fins are usually soldered in the place to form the flow channels. A cover lid is also used in this design to form the flow passage. Different fins have been used to improve the heat transfer performance such as straight, herringbone, wavy and perforated.
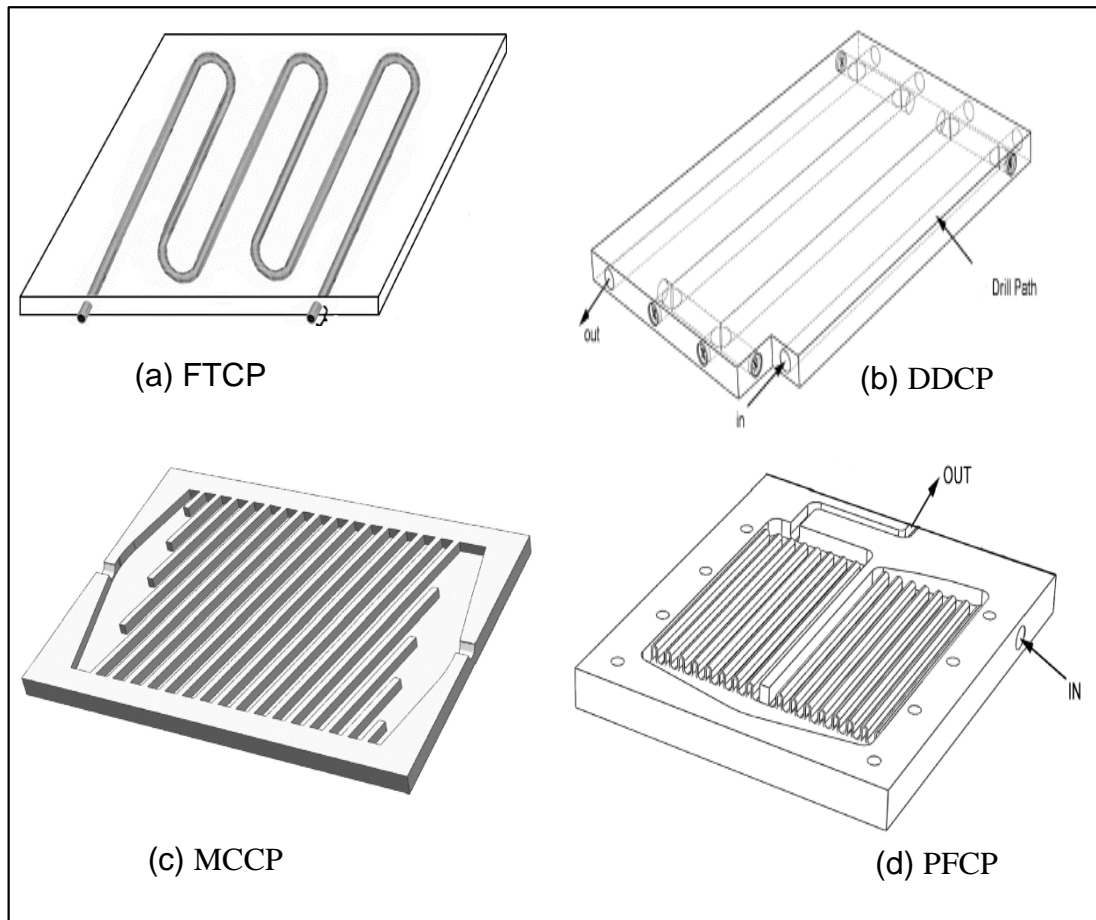
Figure 2.7 Different type of cold plate designs used in DCLC electronics [73].

The CPU temperature is a function of thermal resistance between the active semiconductors of the CPU and the water coolant that passing through the cold plate to take the heat away. The thermal resistance is formed of four layers as shown in figure 2.8: thermal interface material (TIM) between the semiconductors and the CPU lid, the lid material, the TIM between the lid and the cold plate, and the cold plate material. The thermal resistance of the TIM depends on the thermal conductivity, thickness, and the surface area. Ouchi et al. [74] investigated the effect of using three different TIM types which are silver based paste, Taica λgel DP-100, and SCID sheet. They found that the CPU's temperature is higher for the case of using silver based paste compared with the other two TIM for four CPU stress levels 25%, 50%, 75% and 100% as shown in figure 2.9. Schultz et al. [6] modified the traditional designs by attaching the cold plate directly over the semiconductors using

33

liquid metal thermal interface material (LMTIM). Hence in this design, they reduced the thermal resistance formed by the CPU lid and one of the TIMs. The results showed a reduction in temperature of about 6℃ compared to the traditional design of the same type of the CPU under the same IT stress level.

Cold plate

TIM between CPU lid and cold plate

CPU Lid

TIM between the semiconductors and the lid

Chip module

Support frame
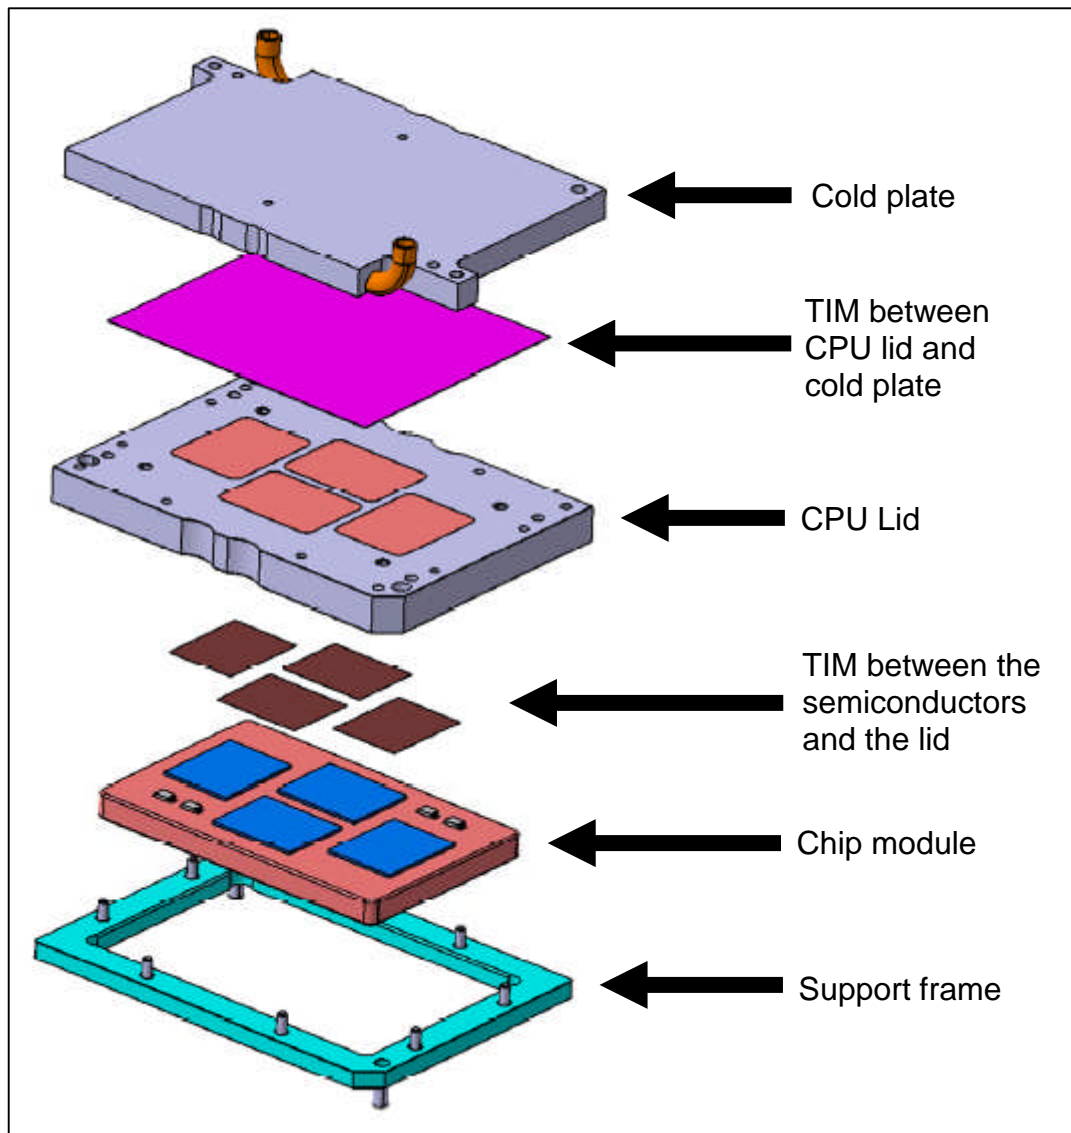
Figure 2.8 Assembly of the chip and the cold plate [75].
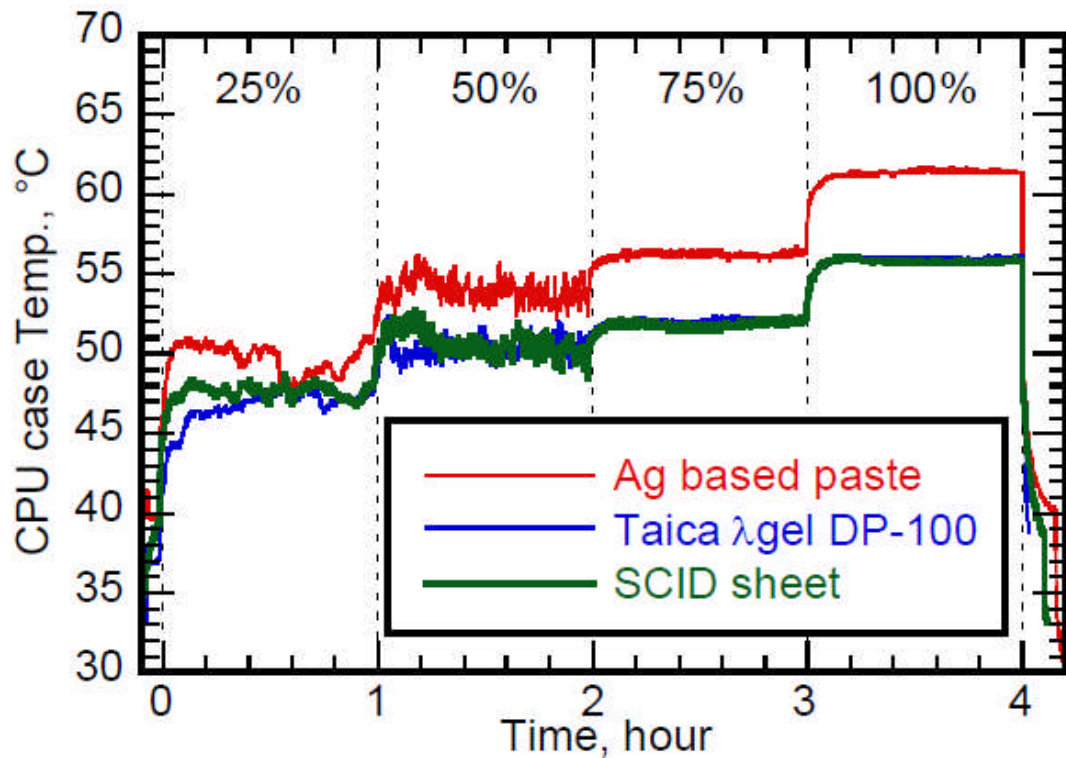
Figure 2.9 CPU temperature comparison for three different types of TIM [74].

The studies of enhancing the cooling performance at the server level are also extended beyond the limits of enhancing the cold plates and the thermal resistance between the chip and the coolant. Ouchi et al. [74] developed a plug-in connector cold plate as shown in figure 2.10. The design aimed to standardise the cold plates and simplify retrofitting new servers and changing the rack configuration. The design evolves attaching a heat spreader to the CPU where the cold plate can be inserted into the slot of the heat spreader secured by a latching mechanism. Kheirabadi and Groulx [76] introduced a new design, where the conventional liquid cooling arrangement is replaced with a liquid/liquid thermal contact heat exchanger as shown in figure 2.11. The main purpose of the design is to increase the reliability by removing the leakage risks. However, using an extra heat exchanger in the loop added an additional thermal interface resistance and reduced the thermal performance. They have also investigated different types of heat exchangers in the design and concluded that using parallel plate types offers less pumping power consumption and acceptable thermal performance.

35

Before insertion

After insertion
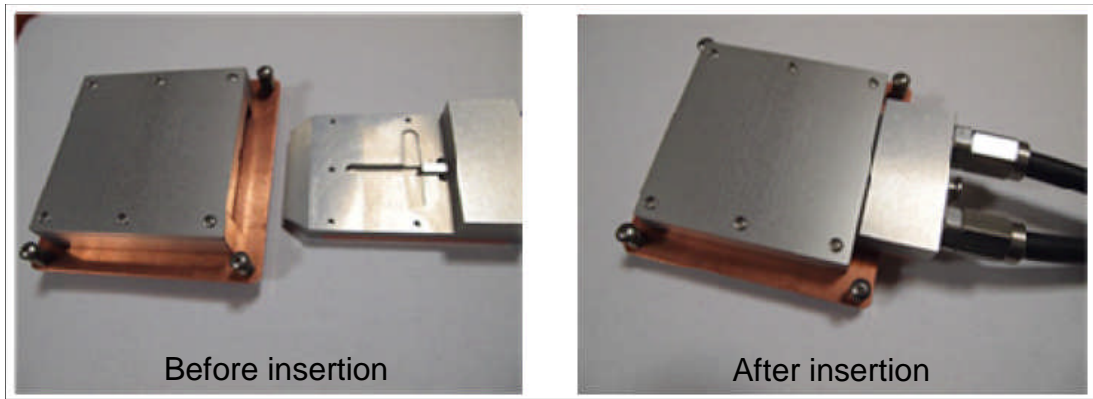
Figure 2.10 Plug-in cold plate connector [74].
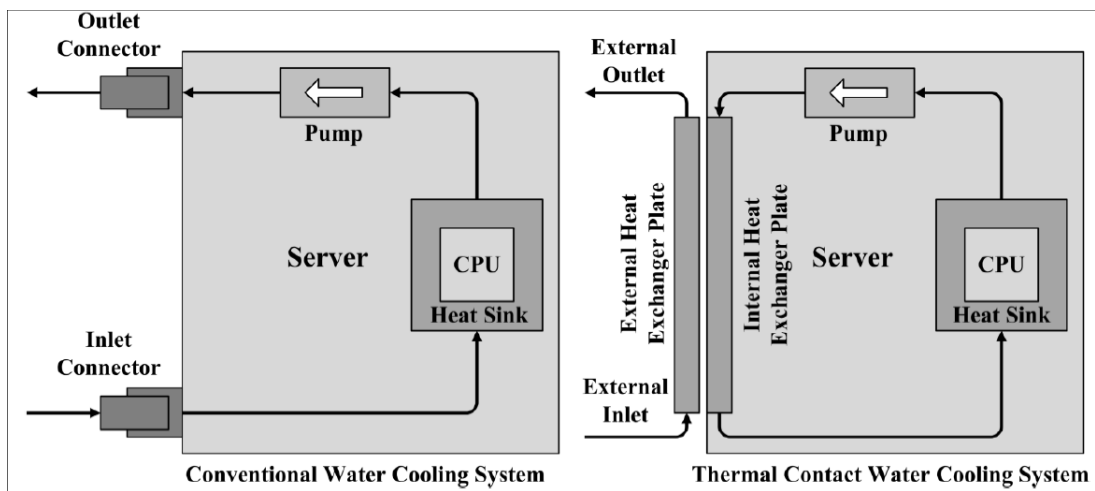
Figure 2.11 Schematic of the thermal contact water cooling system proposed by Kheirabadi and Groulx [76].

Ramakrishnan et al. [77] designed and built a bench level experimental set up to characterise the CoolIT System [78] cold plate. The characterisation study was obtained by calculating the thermal resistance of the cold plate for different flow rates and applied power. The results showed that increasing the flow rate gives lower thermal resistance at the expense of higher pressure drop and pumping power. The input power was found to not have any effect on the thermal resistance of the cold plate.

2.4.1.2 Coolant inlet temperature

The impact of liquid coolant temperature on the thermal and power consumption management of electronics have been studied in the literature at the server level. Exposing the servers' hardware to higher inlet liquid

temperatures leads to improvement in economiser efficiency and reduces the overall cooling power consumption. On the other hand, increasing the inlet water temperature is accompanied with a greater operating temperature of the processor and potential reliability reduction of the operation. Moreover, increasing the processor temperature increases the subthreshold leakage which increases the static power consumption and weakens the performance of the processor [79].

Zimmermann et al. [80, 81] investigated the energy and exergy efficiencies of the Aquasar supercomputer for air and water cooled prototypes. The comparison results between the water cooled and air cooled servers showed that the server can be efficiently cooled with water due to the low thermal resistance of water compared to air. The server can be cooled with water temperatures as high as 60℃ with the temperature difference between the inlet and the outlet of 15℃ whereas a temperature difference of 35℃ is required for the server when it is cooled by air with inlet air temperatures of 23℃. They found that switching to liquid cooled servers increases the exergy efficiency by up to 34% where about 80% of the recaptured heat can be used when the inlet temperature to the server is 60℃.

Addagatla et al. [82] carried out several experiments on a novel design of hybrid liquid/air cooled 2OU (OpenU) web server to investigate the effect of inlet water temperature on the thermal performance and power consumption. In this design, the CPUs are cooled by direct liquid cooling while the remaining components of the server are cooled by air guided by special ducts and driven by the server's fans and a heat exchanger inside the server. The studied range of temperatures was between 27.5 to 45℃ (2.5℃ increments). The study was extended to cover different CPU stress levels: idle, 40%, 60%, 80%,100% and a single experiment includes memory exercise with full CPUs load. The server is cooled by air dry cooler which is controlled by fan speed. The results showed that it is possible to operate the server at the higher inlet temperature range of 45℃. The cooling power consumption of the servers depends on the set inlet temperature and ranges from between 4.92 to 7.86W which results on PUE of 1.54 to 2.4. The fan speeds increased based on the power level of

the server and the water inlet temperature while the water flow rate is kept constant for all the experiments.

Druzhinin et al. [83, 84] studied the impact of inlet water temperature on the power consumption and computational performance of a RSC Tornado supercomputer. The range of the inlet water temperature studied is between 19 and 65℃. They found that the computation efficiency reduced from 2.72 $GFLOPS/W$ at 19℃ to 2.44 $GFLOPS/W$ at 65℃. However, the power consumption of the server increased from 365 to 398W for this increase in the inlet coolant temperature. The study showed that the operating inlet temperature of the coolant at 55℃ gives the best operational conditions as it is estimated that the heat captured can be reused if it is in the range between 40 and 60℃ which also offers free cooling in most of the climate conditions.

Garraghan et al. [85] predicted a unified power consumption model for a Sun Fire V20z server. The novelty of the model falls into its holistic characteristics of capturing the effect of software, hardware, and cooling techniques. The model provides sub-systems power consumption for various utilisation levels. The model was validated with numerous experiments on air cooled and liquid cooled servers. Their results showed that the power profiles are different for variable utilisation levels, server architecture and cooling approach. For this, the relationship between the server power, performance of the application, CPUs temperature and fan speed is not always linear and is also a function of load level and cooling technique.

Sahini et al. [79] studied the impact of high inlet coolant temperature on the static power losses at the server level. The study was achieved experimentally on an Enterprise-class server which offers the opportunity to be fully air or liquid CPU cooled. They found that the CPUs temperature in direct liquid cooling are lower by 20℃ compared to air cooled server. Their results showed that increasing the inlet water temperature from 25 to 50℃ increases the CPUs average temperature by about 21℃ and the power consumption of the server by about 4%.

38

2.4.2 Rack level studies

Direct liquid cooling systems, with their high thermal and energy performance, are of growing interest to the data centre industry. The primary intention and application of existing studies is to obtain maximum advantage of the predominant liquid coolant. Common solutions to optimise the thermal performance of the rack is to supply higher inlet coolant temperature for the purpose of maximising the economisation by reducing the cooling power consumption and increasing the heat reuse.

Zeighami et al [86] used a simplified model to calculate the heat recovery of hybrid liquid/air cooled rack level. The servers used are Cisco C220 M3 CPU liquid cooled and the remaining of the components are cooled by air. The cold plates are provided by Asetek RackCDU D2C^{TM} which have a fixed speed liquid pump. The model is based on the conservation of energy equation with the assumption that the coolant water temperature is the same as the components temperature. The recovered heat was found to depend on the input power, water temperature, and air temperatures provided to the rack as shown in equation 2.18. Three simple parameters were found to dominate the behaviour of the hybrid cooled rack which are: the effectiveness of the captured heat, static power losses due to the current leakage, and the loss of conductance power extended over a wide range of stress levels and inlet water temperatures.

$$P_{recovered} = K_{HR}(P_{server} - P_\circ) - C_L(\bar{T}_{water} - T_{air}) \qquad\qquad 2.18$$

Where, $P_{recovered}$ is the recovered energy, P_{server} is the measured server power, $P_\circ$ is the static base power loss, $\bar{T}_{water}$ is the average water temperature between the inlet and the outlet, T_{air} is the measured inlet air temperature, C_L is the power loss conductance constant and K_{HR} is the power recovery effectivness.

Li et al. [87] proposed an algorithm, which is called SmartPlace, that arranges liquid cooling servers based on the workload distribution to reduce the total power consumption of the data centre. The proposed model is also taking into account the optimum distribution of the load by putting some servers into idle

operation for the goal of minimising the total power consumed by the IT and cooling. The SmartPlace algorithm was applied on two data centres of 1280 and 10240 servers capacity and the results showed more than 26.7% power saving was achieved a high dynamic response performance.

Ovaska et al. [88, 89] compared the effect of inlet temperature on the Norwegian Stallo supercomputer under fully air cooled and CPUs liquid servers. The LINPACK benchmark exercise was run on the servers during the experiments. Three inlet temperatures of 20, 25 and 30℃ were chosen for the air and liquid cooled servers and the results were presented in terms of rack power consumption and computational performance. They found that the power required to perform 1 Gigaflop/s of computing is about 1.04 to 1.11W for the air cooled rack, while it is only 0.93 to 0.94W for the water cold rack.

2.4.2.1 Flow distribution and pumping configuration

The coolant flow rate is also a potential subject of investigation since higher flow rates mean higher pressure drop and pumping power and lower flow rates reduces the reliability and risk of shutting down due to the increase in CPU temperatures resulting from insufficient flows. However, the design objectives of the rack systems fall into two aims: providing sufficient flow and distributing the flow evenly in the servers to maintain the temperatures inside the electronics at the designed level.

Kelkar and Patankar [90] introduced the technique of Flow Network Modelling (FNM) using MacroFlow software to analyse the flow in a distributed manifold as shown in figure 2.12. The flow results showed to be non-uniform due to the major and minor losses in the manifolds. They found that increasing the cross-sectional area from 7/16 inch to 7/8 inch have reduced the flow rate difference between cross-1 and cross-7 from 0.6gal/min to 0.08gal/min, respectively.
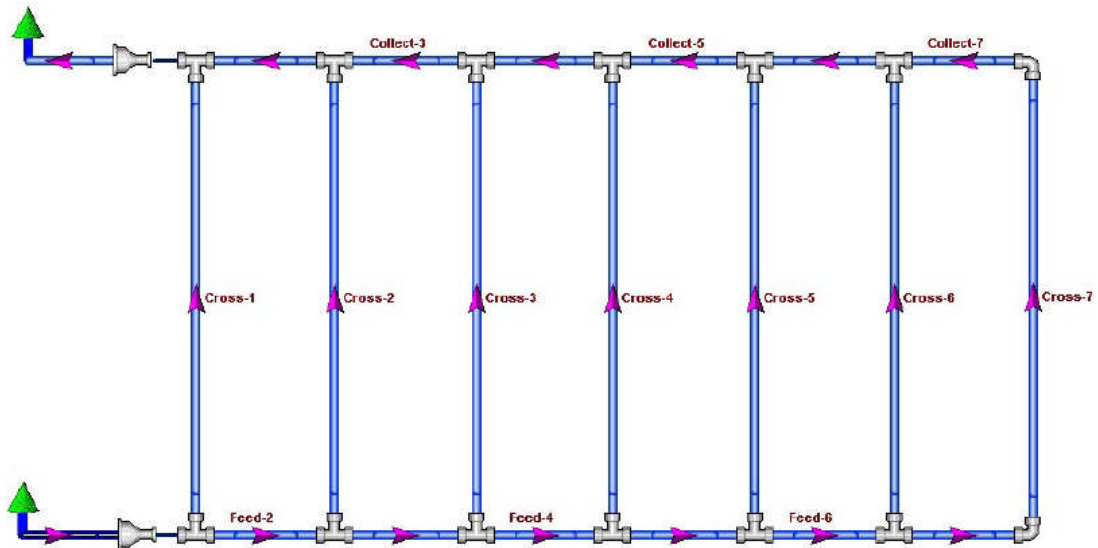
Figure 2.12 Network model of the manifold distribution [90]

Ellsworth et al [91] carried out a flow analysis on a rack of IBM 775 Supercomputer servers. The rack contains 12 servers, disassembles, 2 bulk power assemblies, and a rear door heat exchanger as shown in figure 2.13. Four pumps and liquid/liquid heat exchangers are positioned in the bottom of the rack and connected in parallel to increase the redundancy. The study was aimed to assure provision of minimum flow rate required for all the components in the rack and minimise the pumping power consumption for a different range of rack configurations. For this, the MacroFlow software was used to analyse the flow and model the flow network. The model was validated against the experimental data with an accuracy of 5%. The results showed that the flow in each branch is higher than its minimum requirements for their design. However, there is a maldistribution in the flow of the identical branches which theoretically have the same pressure drop as shown in figure 2.14.

Alkharabsheh et al [92] presented an experimental and numerical pressure drop analysis study of a rack level containing servers which are cooled by DCLC cold plates from CoolIT Systems. The pressure drop in the direct liquid cooled system is directly linked to the power consumption in the rack due to pumping. The main aim of their study is to evaluate the pressure losses and flow distributions in the whole rack and through the individual servers. In addition to that, they evaluated the accuracy of FNM in modelling the flow

41

network in the direct liquid cooling system. The rack design contains mainly three components: Coolant Distribution Module (CDM), supply and return manifold module, and server cold plate module. The experiments showed a flow maldistribution through the identical server modules and they attributed this to the manufacturing limitations of the microchannel cold plates. The results revealed that the server modules have the highest portion of the pressure drop in the rack of about 59% among the CDM and the manifolds. However, 69% of the pressure drop in the server's module is caused by the corrugated tubes, fitting, and the stäubli valves. The reason for that is that the server's modules are designed to be flexible and provide a reliable quick connection which results in higher pressure drop. The FNM mode is validated with the experimental data for 15 servers in the rack and the CDM module, which resulted in an average accuracy of 2.5% and 3.8%, respectively.
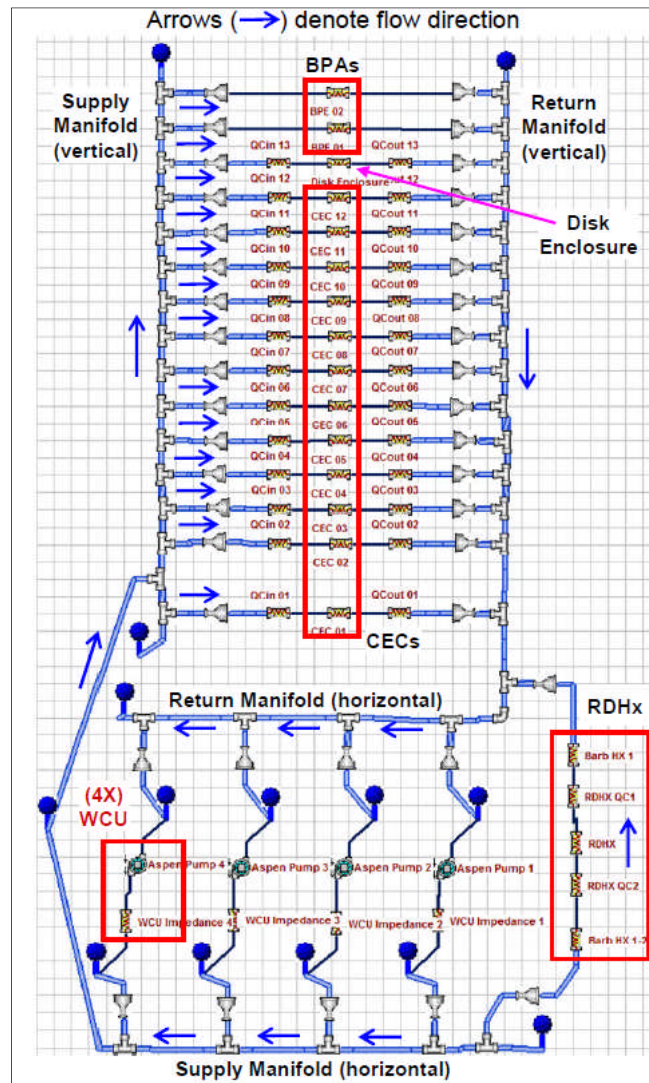
Figure 2.13 Flow network representation of 775 IBM rack system [91].

Sahini et al. [93] studied the effect of inlet water temperature on the power consumption and component temperatures in a mini rack of OpenU (2OU) web servers as shown in figure 2.15. The study also provided a comparison between the pumping system of central and distribution. The centralised pumping is where the flow is driven by external pumps while the distributed pumping is where the flow is driven by the small pumps which are integrated on the cold plates. The pump types in both configurations are of constant pumping speed. The range of studied water temperature inlet is between 25 and 45℃ following ASHRAE liquid cooling envelope W4. Their results showed that the CPUs temperature is lower for centralised pumping compared to the distributed pumping.
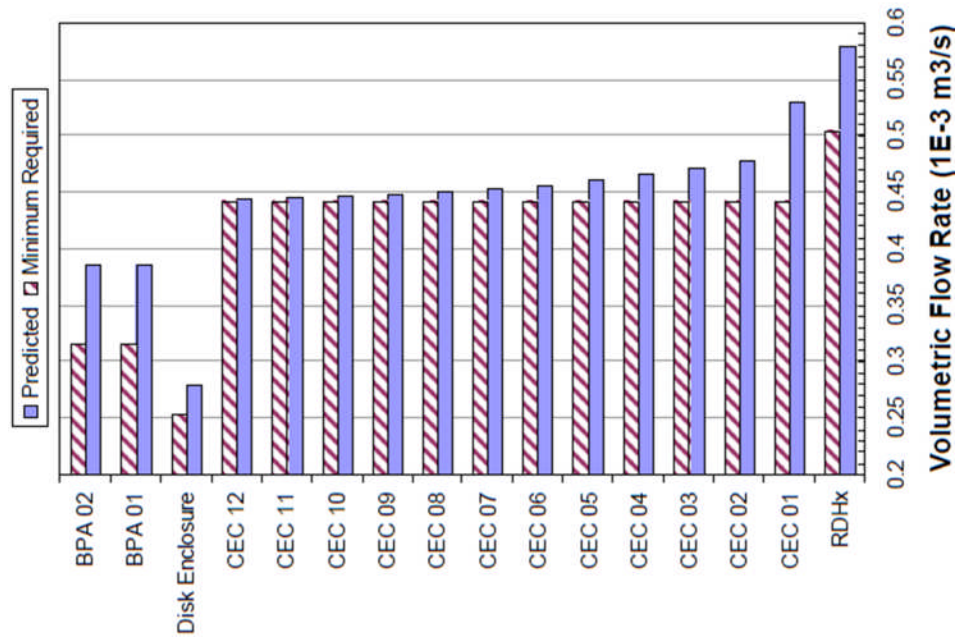
Figure 2.14 Flow rate distribution of the 775 IBM rack predicted using MacroFlow software [91].
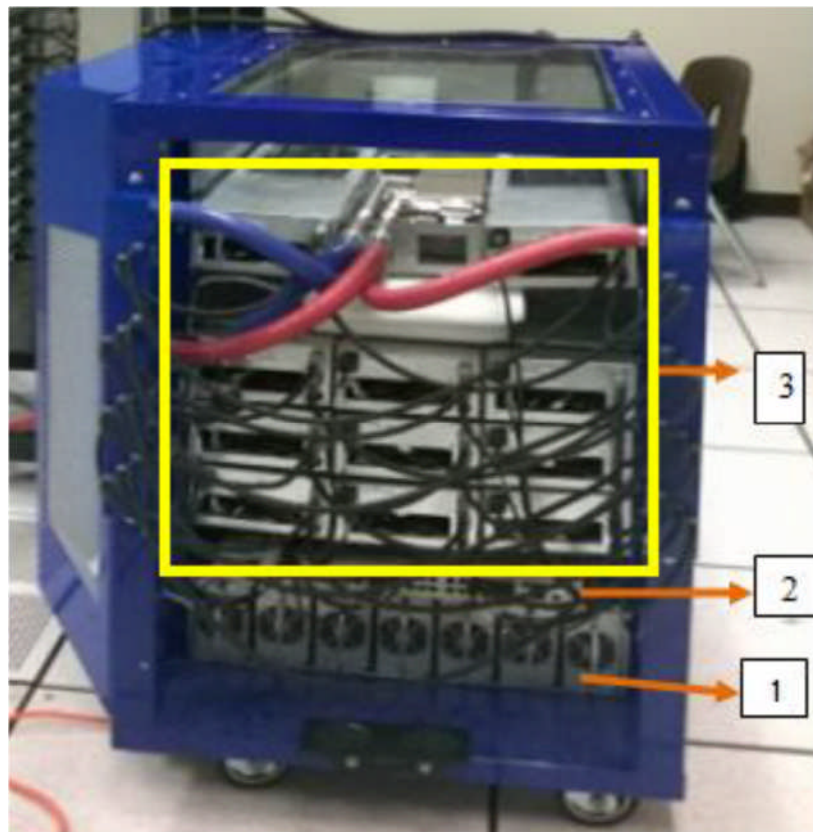
Figure 2.15 Rack system studied by [93]. (1) Power shelf, (2) network switch, and (3) 2OU servers.

2.4.3 Facility level studies

The increasing amount of heat rejected by data centres has increased the need for implementing and enhancing the performance of different heat rejection systems. The heat rejection systems of data centres vary depending on the geographical location, size of the data centre and available resources. However, the number of studies that detail the facility level is limited compared with studies for other data centre level components.

A chiller-less data centre based on ambient air cooling that utilises both air and liquid cooling at the rack was patented by Chu et al. in 2010 [94]. In this patent, they presented the design of IBM servers in a rack which are cooled using a combination of direct liquid cooling and air cooling. Iyengar et al [95] and David et al [96] presented more details about the IBM chiller-less design termed as "dual enclosure liquid cooling" which is shown in figure 2.16. The dual enclosure liquid cooling means that the cooling system consists of two liquid loops (internal and external loops) connected thermally through a liquid/liquid heat exchanger which is called the buffer unit. The internal loop is carrying the heat indirectly and directly from the servers. In the indirect way, part of the heat is transferred to the internal loop through an air/liquid heat exchanger within the rack enclosure which they named it "sidecar heat exchanger". The water of the internal loop then enters the servers through the liquid cooled components to be in direct contact with the heat generation sources which are mainly the CPUs and the RAM. The internal loop then exchanges the heat to the external loop through the buffer unit. The external loop then carries the heat to a dry cooler outside of the building which uses air/liquid heat exchangers assisted by fans.
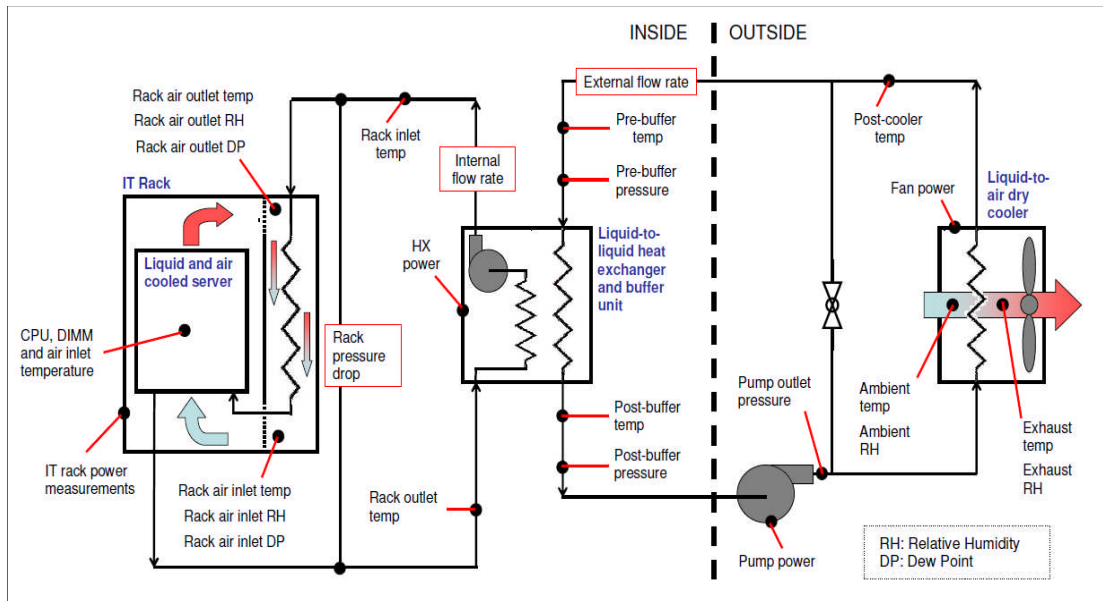
Figure 2.16 Dual enclosure liquid cooling system of IBM [96].

A number of studies were carried out by the IBM Systems and Technology research group on the previous design to improve the performance and characterise the heat exchange points and the operation under different ambient conditions. David et al. [96] investigated the impact of the operation conditions of flow rates, heat exchangers arrangements, the addition of glycol, and weather conditions on the thermal performance of the design. They found that adding propylene glycol to the loop with a mass ratio of 20% increases the temperature inlet to the servers by about 7 to 10°C due to the poorer thermal conductivity compared to water. This indicates that using propylene glycol requires increasing the loop flow rate to compensate for the deterioration in the heat transfer which increases the pumping power as the viscosity is also increasing with adding the glycol. Iyengar et al. [97] proposed a simple thermodynamic model for the design validated by long run experiments. The model was then used to investigate the feasibility of using the design for different cities in the US. The results showed a significant energy saving using chiller-less data centre cooling design which could reduce the cooling energy consumption to about 3.5% of the IT power at the data centre level. Parida et al. [98] highlighted the possibility of using a single loop in the design instead of a dual loop. They found that the power consumption for the single loop is higher than the one for the dual loop under the same

46

ambient and operation conditions due to the low performance of the antifreeze. They concluded that using single loop could have a significant saving in the areas of high ambient temperatures where adding antifreeze is unnecessary. However, the raised concerns of water quality (e.g corrosion inhibitors) make it necessary to use dual loop cooling system where high water quality can be used in the internal loop and lower water quality can be used in the external water loop.

Based on the previous design, Parida et al. [99] developed an energy efficient servo algorithm to control the power consumption of the cooling system to a minimum based on the targeted supplied temperature to the rack and the ambient conditions. The algorithm was designed into three zones depending on the inlet temperature to the rack and the targeted set temperature as shown in figure 2.17. They defined zone 1 when the inlet temperature to the rack is below the targeted minimum temperature where the recirculation valve is used to control the temperature while zone 2 when the inlet temperature to the rack is between the minimum and maximum temperature set value. Zone 3 is defined as the situation where the inlet temperature to the rack becomes higher than the targeted set temperature. Zone 1 and 2 offer high energy saving where no work is required by the pump and fans. While at zone 3, the dry cooler fans and pump speeds need to be increased based on the difference between the inlet temperatures to the rack and the target temperature. The proposed algorithm offered more than 25% energy saving at the data centre level and more than 90% saving compared with the refrigerant based data centre cooled system.

Meyer et al. [100] presented the feasibility of using hot water cooling for energy reuse in the iDataCool HPC system, which was jointly developed by the University of Regensburg and the IBM Research and Development Lab Böblingen. The racks are based on IBM's iDataPlex platform which was customised to be hot water cooled at an inlet temperature of $70℃$. To investigate the energy reuse, an adsorption chiller was coupled with the system which operates at high performance at these temperatures as shown in figure 2.18. In the absorption chiller, the refrigerant is thermally compressed

using liquid refrigerant/sorbent solution and the harvested heat from the data centre instead of using the mechanical compressor as in the conventional chillers. The cooling effect is initiated in the evaporator due to the evaporation of the refrigerant at a very low pressure then to the absorber where the vaporized refrigerant is absorbed to dilute the solution (usually LiBr or ammonia/water). The solution is then pumped to the generator where the solution is regenerated using the driving heat that collected from the data centre. The generated chilled water from the absorption chiller can be supplied to cool the other parts of the data centre. The results showed that about 50% of the energy can be recovered to provide chilled water. This is particularly useful in the cases where the hot water is not needed for heating where it can be used to provide chilled water. They also indicated that there was no negative effect on the IT after leaving the servers working at this range of temperatures for more than one year.
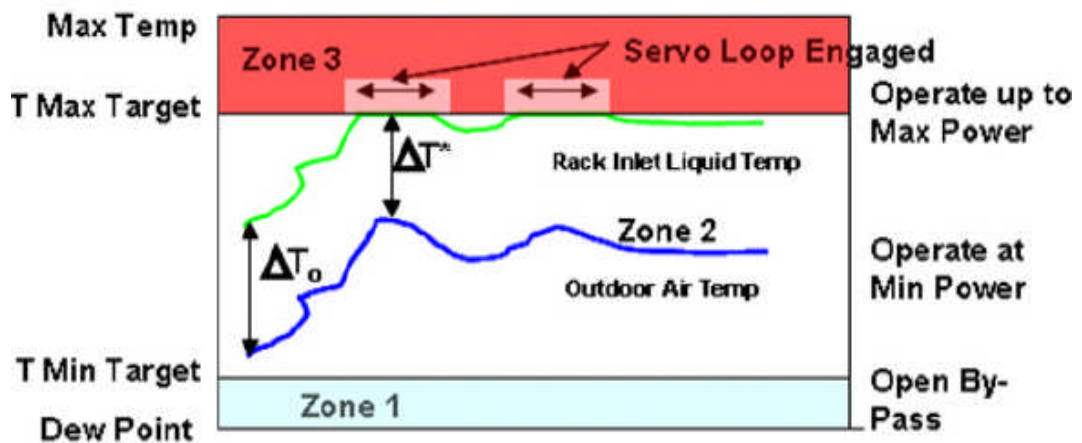
Figure 2.17 graphical representation of the servo control algorithm proposed by Parida et al [99].

48
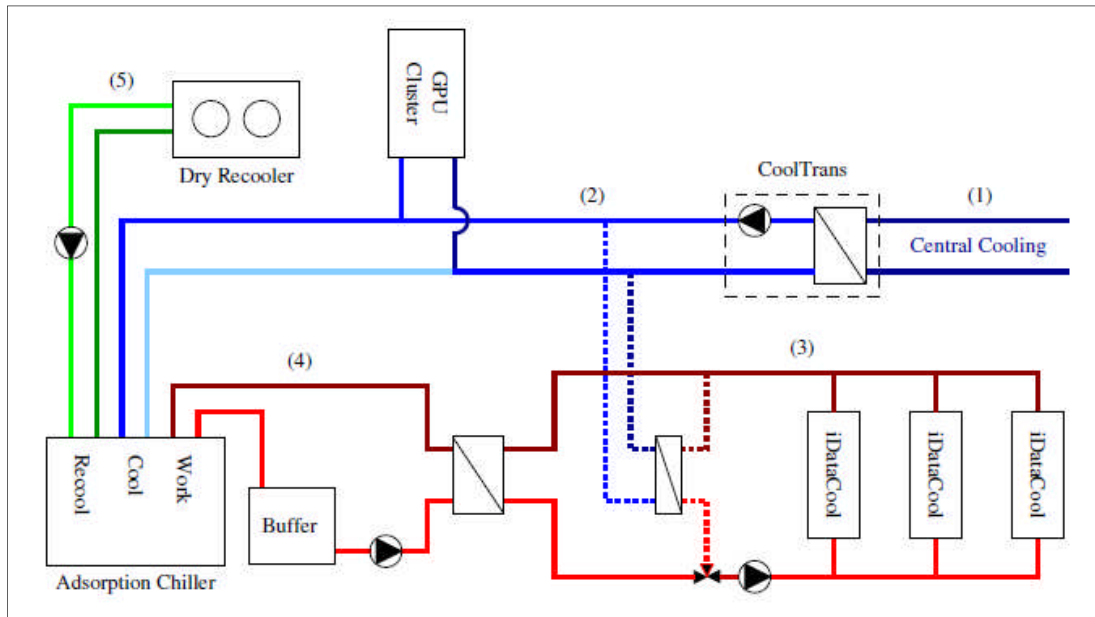
Figure 2.18 liquid cooled data centre proposed by Meyer et al.[100]. (1) central cooling system, (2) primary cooling loop, (3) cooling loop of the racks, (4) driving loop (5) re-cooling loop.

2.5 Overview of data centre liquid cooled designs

As described previously, the heat generated in data centres is due to the operation of the microscale components inside the servers. The larger portion of this heat is generated mainly by the CPUs. However, depending on the server type, the other portion of the heat is generated by RAM, Hard disks, power supply units, and the other auxiliary hardware components of the server. Some of these components are easily cooled by direct water contact and others are difficult and add more complexity to the design. Therefore, cooling these servers by direct liquid contact techniques can be achieved for some components of the server. However, in all cases, air cooling is still required to cool some parts of the servers which are difficult/impossible to be cooled by the direct contact liquid cooling technique.

As the CPUs are the highest heat generation part of the server, these components are generally cooled using cold plate techniques while the other components of the server are air cooled as shown in figure 2.19 [79, 88, 89]. Although in some designs the RAM is also included with the direct liquid cooled technique as shown in figure 2.20 [6, 86, 101]. However, other designs

have been extended to cover more components of the server by the direct liquid cooling technologies such as the IBM 775 supercomputer water cooled design as shown in figure 2.21 [67, 75, 80].

Using water to cool the components of the server directly improves the thermal performance of the components and helps to increase the economisation efficiency. However, there are many challenges associated with using the direct liquid cooling technique to cool the other elements of the server apart from the CPUs and RAM. Goth et al. [75] mentioned that several issues raised with attaching cold plates to the most of the server components shown in figure 2.21. These challenges are first associated with the high weight and cost of the server, secondly, the issue of the complex serviceability of these servers, thirdly, the difficulties associated with replacing or upgrading any item inside the server and finally, the issue of thermal protection in this type of design is more crucial compared with air cooled server where special design is necessarily taken into account to keep the server safe in case the flow is clogged in one of the components.
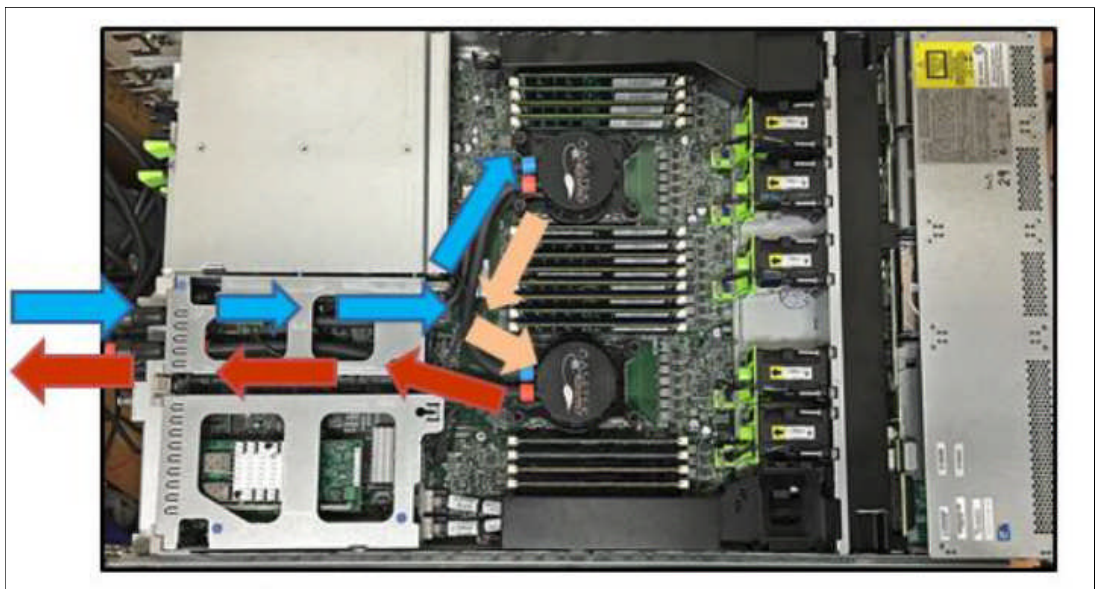
Figure 2.19 1U Enterprise class server retrofitted with a direct liquid cooled technique from Asetek^{TM} [79].

In all the design cases of the direct liquid cooled server, part of the heat generated by the servers will be transferred by air. However, the amount of this heat depends on the type of the server and which components are

covered by the direct liquid cooled design. For this, it is necessary to capture this heat and take it away from the IT enclosure and the IT room. This can be achieved in two ways: using CRAC/CRAH unit in the IT room or utilising the rear door heat exchanger technique (RDHx) or sidecar heat exchanger (SCHx). While the CRAC/CRAH is a very well-known technique to data centre air cooling, RDHx and SCHx are the key development steps in data centre cooling that brought the water closer to the rack before invading the IT itself by the direct liquid cooling.

The RDHx technique has been under extensive investigations by many researchers to reduce the power consumption in air cooled data centres [102-109]. The RDHx is fin and tubes liquid/air heat exchanger type that mounted in the back of the rack where chilled water passes through the tubes of the heat exchanger to cool the air that passes through the fin surfaces as shown in figure 2.22 [102]. The concept of the SCHx is very similar to the RDHx apart from the position where the SCHx is positioned to the side of the rack and contained with an enclosure inside the rack to guide the air flow through the fins surfaces as shown in figure 2.23.
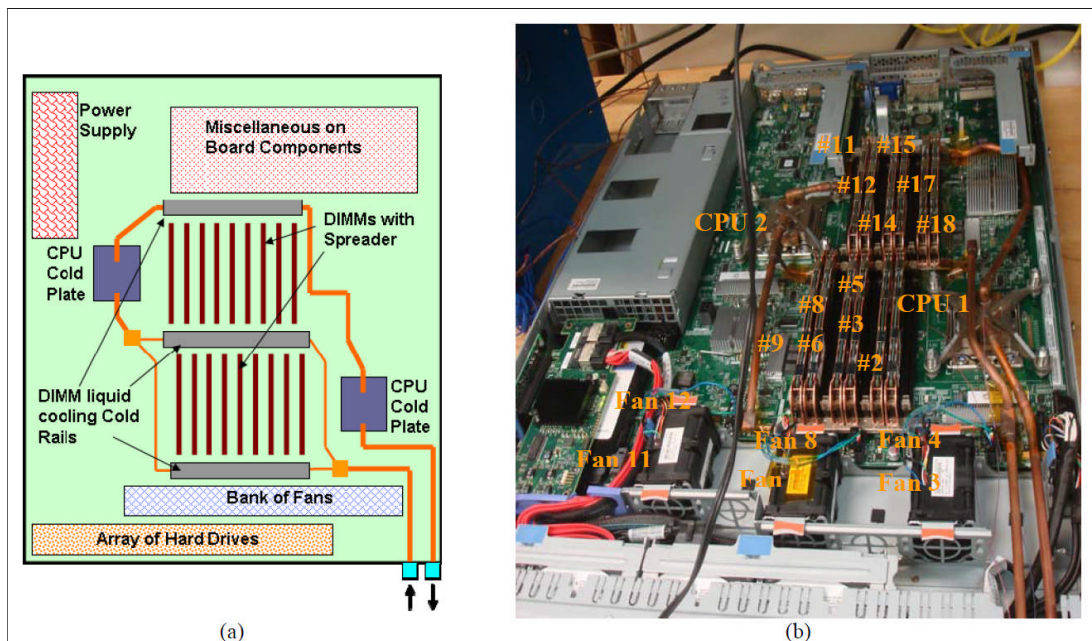
Figure 2.20 IBM system volume server uses direct liquid cooling for the CPUs and RAMs, (a) schematic of the server and (b) is the server picture [6].

The RDHx can work in two different scenarios of air flow: passive and active. In the passive operation, the air flow is generated by the small fans inside the servers of the rack. Under the active mode, the air flow is also generated by the servers fans but with the assistance of a number of big fans mounted on the RDHx. Almoli et al [109] carried out a three dimensional CFD analysis to compare the performance of the RDHx in the active and passive operation and found that the active RDHx offers 2% improvement in the thermal performance compared to the passive RDHx.

The RDHx is originally designed for fully air cooled racks to take 40 to 90% of the heat away while the remaining heat is removed by the CRAC/CRAH of the room [109]. Therefore, a large energy saving is accomplished by reducing the number of CRAC/CRAH units in the room. However, in the direct liquid cooled designs, the RDHx and SCHx work on taking the generated heat of the remaining components of the servers which are still air cooled [110].
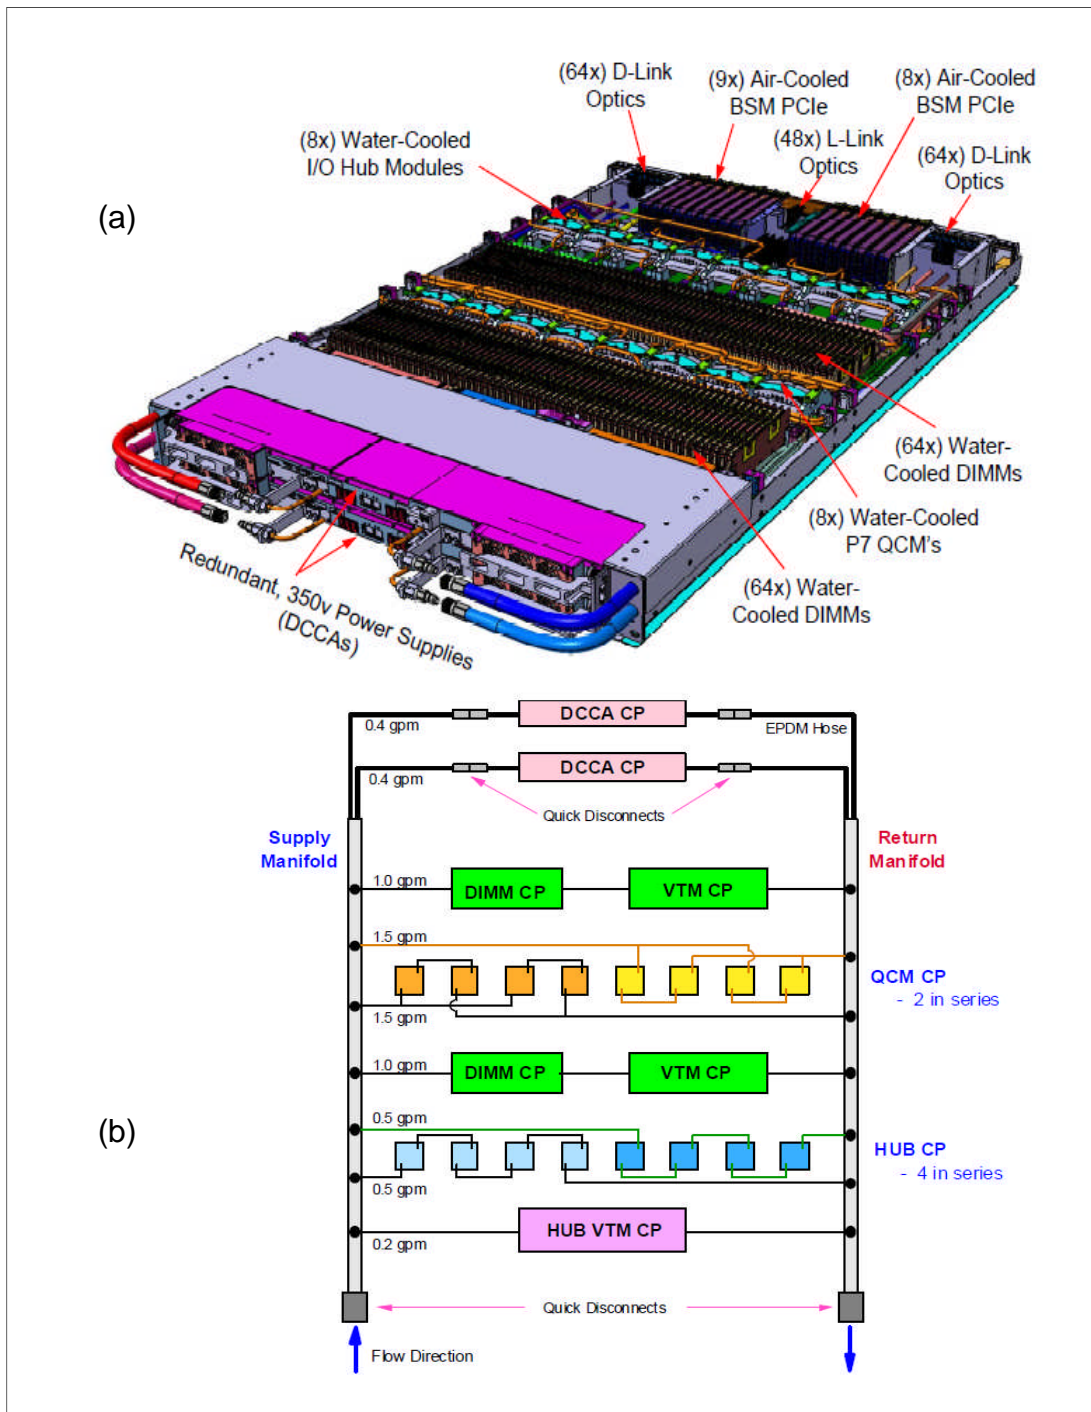
Figure 2.21 Power CEC server, (a) isometric view and (b) server flow diagram [75].

53
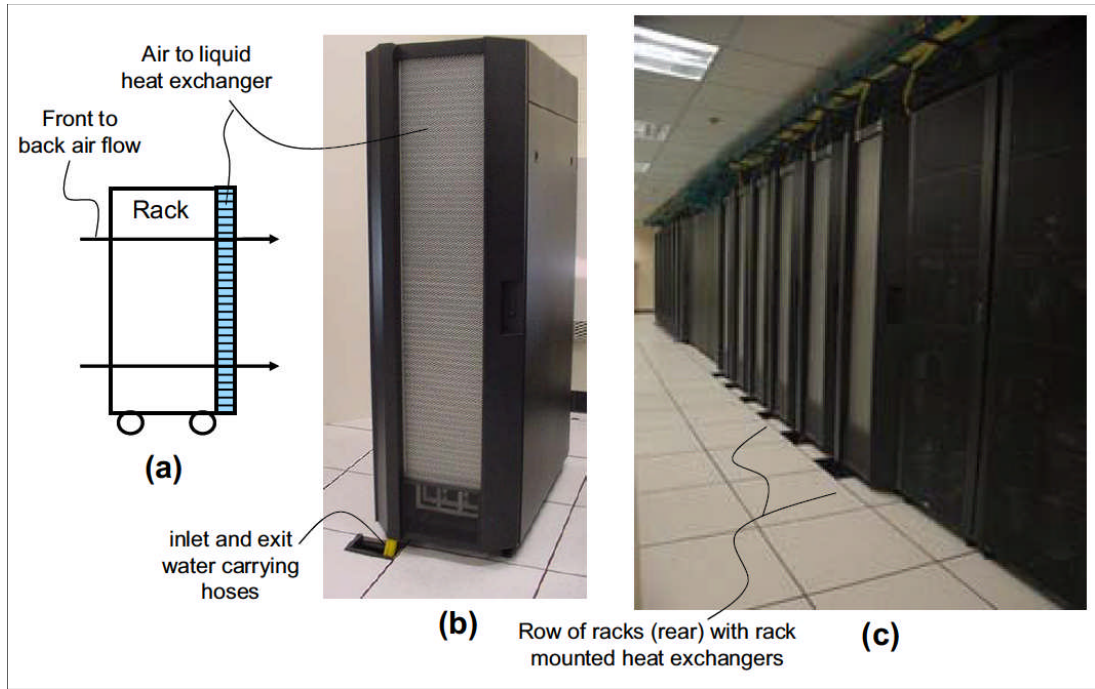
Figure 2.22 Rear door heat exchanger configuration, (a) schematic side view, (b) single rack mounted and (c) data centre configuration [107].
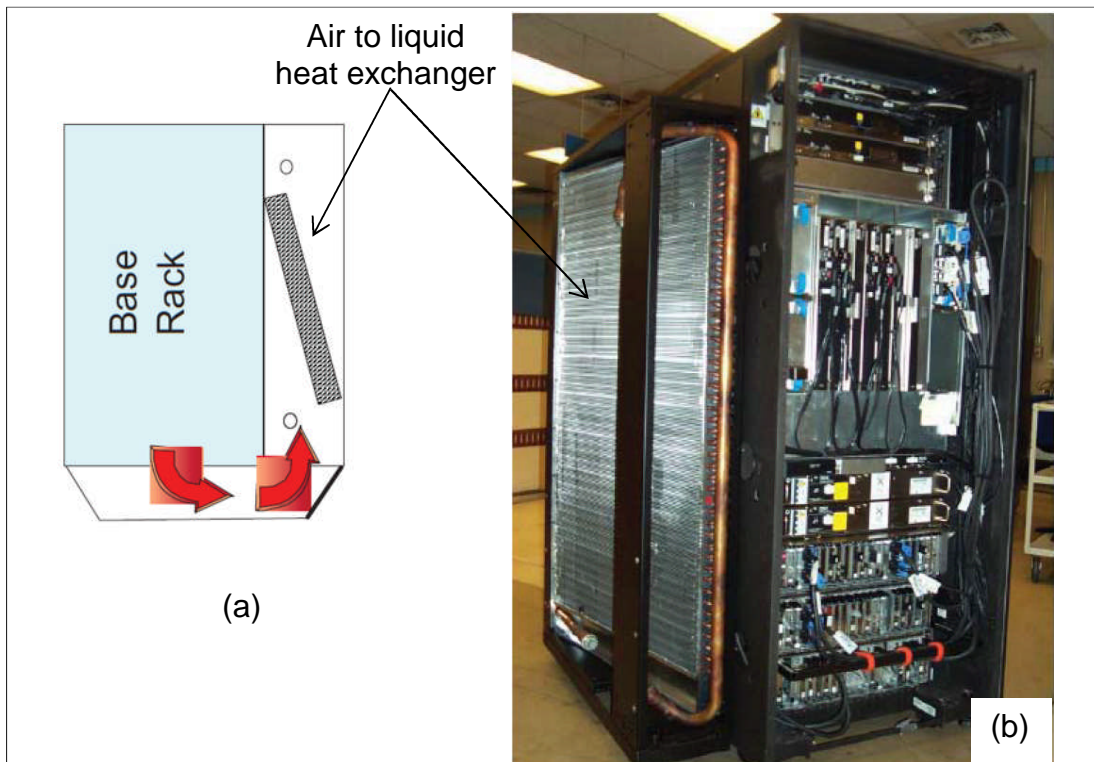
Figure 2.23 Sidecar heat exchanger, (a) schematic top view and (b) single rack design (shown without the cover) [103].

2.6 Summary

This chapter presented a general overview of data centres, energy effectiveness, efficiency metrics and the available cooling techniques. Among the cooling techniques, an emphasis was placed upon the available studies of direct liquid cooling of data centres. These studies were categorised based on the scale that covered by the study into three levels: server level, rack level and facility level. The general aim of these studies was to give a comprehensive picture of energy and thermal management in data centres. The findings of these reviews are summarised as follows:

1. At the server level, studies have focused mainly on two aspects to improve the energy and thermal performance:

    a. Reducing the thermal resistance between the electronic components and the coolant. This was accomplished by either enhancing the available solution (e.g optimising the fin shape and spacing) or proposing a new technology (e.g reducing the thermal resistance layers by using a new thermal paste).

    b. Supplying higher inlet coolant temperature to the server aiming to increase the efficiency of the economiser for lower cooling power consumption, with the potential for higher energy reuse. However, this is accompanied with an increase in the component temperatures (e.g CPUs) which in turn results in lower computational performance, higher power consumption and reduced service life of the server.

2. At the rack level, the literature focuses on three aspects:

    a. Supplying higher temperature to the rack to increase the heat reuse capacity and decrease the cooling power consumption and predict models to capture the recovered heat.

b. Distribute the workload between the servers in the rack and across the data centre.

c. Coolant flow distribution and the flow design within the rack.

3. At the facility level, several studies were carried out on the IBM liquid cooled chiller-less design. These studies have focused on the design of the system, the impact of ambient temperature and different configurations of the design.

Collectively, these studies demonstrated a wide knowledge of energy efficiency and thermal management in a specific level of the data centre infrastructure. There is however, a lack of complete understanding of the holistic management in direct liquid cooled data centre. Moreover, the complexity of the integrated interaction between cooling, workload, power consumption and efficiency has raised many questions for the optimum operational strategy of liquid cooled data centres. Therefore, this thesis will address direct liquid cooled data centres by integrating all the interactions at all levels from chip to facility.

The following points distinguish the present work from those within the wider literature:

1. Only a single study was found to present a design for direct liquid cooled data centres [94]. However, little attention was given to the heat rejection system (dry cooler) and hence a lack of designing detail is available for the complete design methodology. Therefore, this study will propose a novel design of a dry cooler that utilises spray evaporation to increase the heat transfer capacity. Full details will be given from the heat exchanger selection to the control methodology and commissioning to propose full design guidelines for an energy and thermal efficient direct liquid cooled data centre.

2. There is only small guidance about the effect of flow rate in the rack on the component temperatures and a single study was found to investigate the pumping configuration (central and distributed) on the component thermal

performance of a mini web server's rack [93]. The current study will address the effect of rack flow rate for different pumping configurations on the thermal components performance, efficiency and effectiveness of the data centre.

3. While Ellsworth et al [91] and Alkharabsheh et al [92] have highlighted the rack flow maldistribution due to the losses in the manifolds and pressure drop losses in the cold plates, there has been no details about the effect of this issue on the CPU temperature. For that reason, this study will look at the flow distribution in the rack and the resulted CPUs temperature. In addition, this study will extend the study of the flow distribution in the rack to both types of pumping: central and distributed pumping.

4. Existing studies on supplying higher rack temperature inlet were mainly concentrated on the heat reuse and energy saving in the cooling system while few studies have focused on the effect of rack inlet temperature on the holistic performance of the data centre. Hence, this study will holistically investigate the effect of rack inlet temperature for different workloads and calculate the power usage effectiveness and the compute efficiency as well as the comprehensive metric between the two to account for the feasibility.

Chapter 3 Design and construction of hybrid heat rejection system for data centre cooling

3.1 Introduction

The use of direct liquid cooling at the rack coupled with economization could potentially yield significant energy savings and improve the operational performance of the electronics. To investigate the thermal and energy advantages of this concept, a chiller-less, ambient air-cooled data centre with cooling at the rack through a combination of liquid and recirculated air is proposed. The heat dissipated from the IT is rejected to the environment through the use of a closed loop liquid to air (or to sprayed air) heat exchanger (HE). A detailed review and justification of the experimental arrangement has been designed and implemented.

The proposed design of a data centre in this study consists of several levels of systems and subsystems compromising the IT equipment as the primary source of heat and cooling equipment as the heat extraction system. Hence the experimental setup consists of two sides; IT environment side and heat rejection environment side as depicted in figure 3.1.
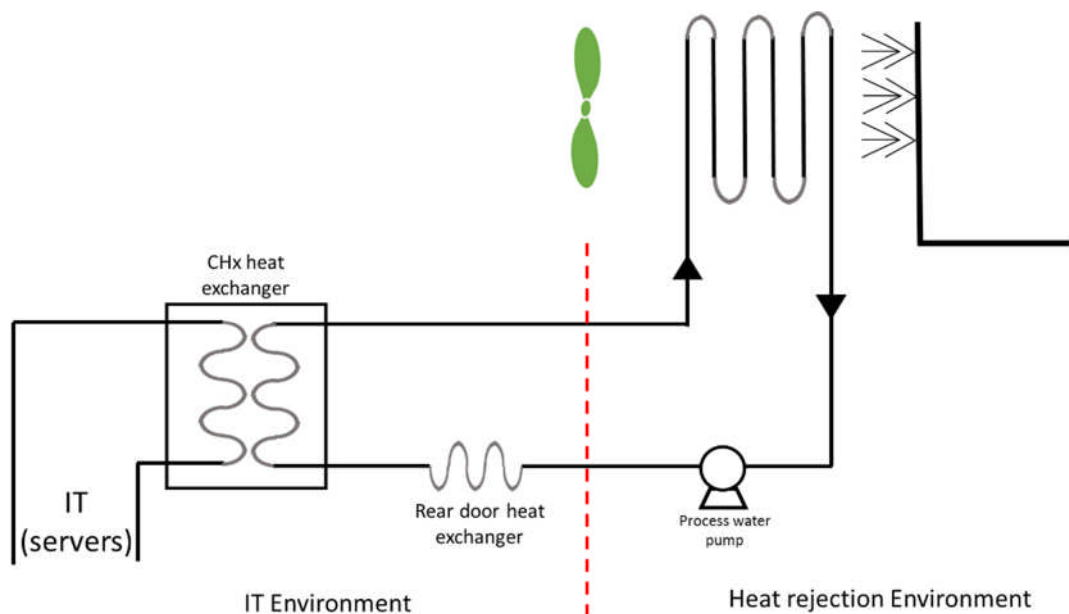
Figure 3.1 simplified sketch of the design.

The water cooled by the air handling unit (AHU) enters the rear door HE and gains some heat from the recirculated air to cool the RAM and PSUs of the servers. After flowing through the rear door, the partially heated water enters the CHx40 HE and circulates back to the AHU HE. The coolant passing through this loop will be called the process water loop or the primary loop. The hot water coming out of the servers is cooled by exchanging its heat to the primary water loop through the CHx40 HE before being pumped back to the servers. The coolant passing through this loop will be called the secondary water loop. The servers' cooling design in the present study accommodates for both, air-cooled and liquid-cooled devices at the server level.

This chapter will describe in details the whole experimental design, starting from the server level to the rack level and finally to the AHU heat rejection system level.

3.2 Servers configuration

Sun Fire V20z servers from circa 2005 [111] are used to represent the IT. Each server consists of 2 CPUs (Central Processing Unit) which are AMD (Advanced Micro devices) Operton 64-bit processors running Debian Ubuntu, a hard disk drive, and 8 DIMMs of installable memory. The first CPU is called the primary (master) CPU and the second is called secondary (slave) CPU which are also called CPU0 and CPU1 respectively. The CPU is a dual core type, 2 GHz frequency and data width of 64-bit and has its own random access memory (RAM). Each RAM bank has four slots of a capacity ranging from 256MB to 8GB. A perspective view of the components, features, and physical specifications are shown in figure 3.2, table 3.1 and table 3.2, respectively.

The thermal design power (TDP) is 85.3 W per CPU, not exceeding a critical die temperature of 71°C. The thermal energy is dissipated through a copper heat sink that sits on the CPU with a thin layer of thermal conductive paste which is spread between the CPU and the heat sinks to reduce the thermal resistance.
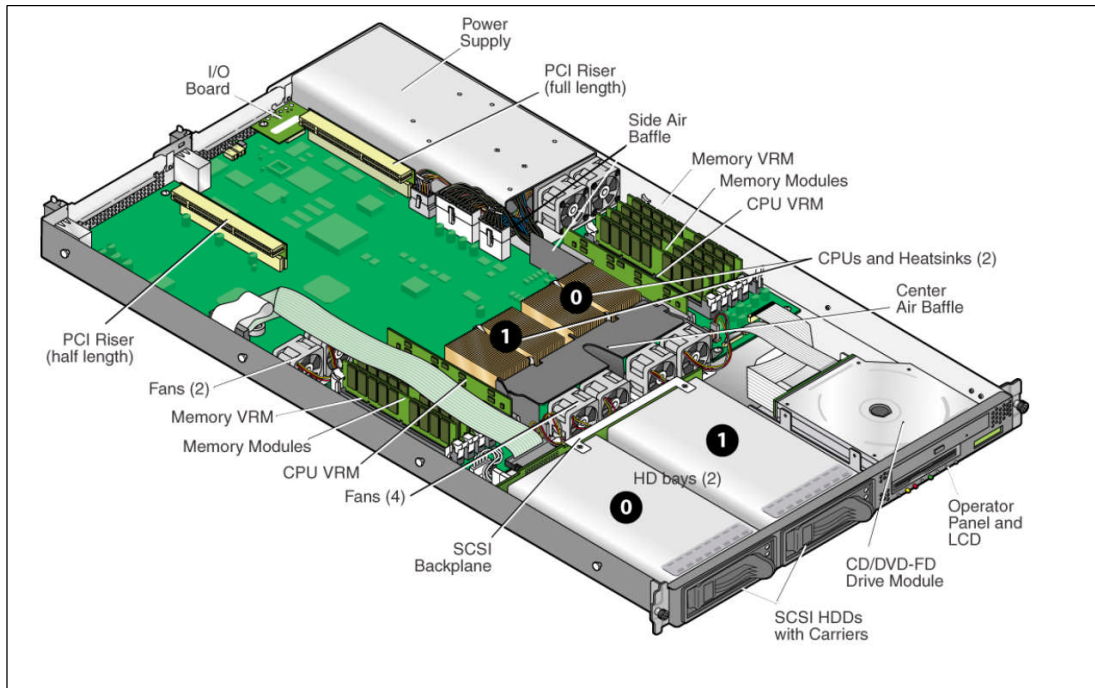
Figure 3.2 Components illustrations of Sun Fire V20z server [111].

Table 3.1 Features of Sun Fire V20z server (2005).

Component	Description
CPU	2 AMD Opteron processors, 64-bit x86 architecture.
RAM	8 slots, 512 MB to 16 GB ECC, registered DDR.
Hard disk drive(s)	1 or 2 U320 SCSI HDD(s).
Network I/O	Dual embedded Gigabit Ethernet.
Other I/O	Internal CD-ROM (or DVD-ROM), Diskette drive, Embedded SVGA Video, Keyboard and Mouse.

Table 3.2 physical specifications of Sun Fire V20z server (2005).

Parameter	Value
Mass	15.88kg
Height	1U (rack unit = 44.45mm)
Width	10U
Depth	16U

To provide the management tool, the service processer (SP) is embedded in the server which is a Linux based operating system that uses Secure Sockets Layer (SSL) encryption to provide the required connection security. The servers are connected to all lab machines via two dedicated 10/100/1000 Mbits/s connections.

The intelligent platform management interface (IPMI) is used to manage and monitor the operational conditions of the server [112]. Through connecting the IPMI to an external server (the head node), it is possible to access the local host of any server in the rack to monitor and log the data of the sensors in the server. The sensors are embedded inside different components of the server which are reporting the operational conditions of these components during the server operation. This includes temperature sensors which measures the ambient temperature inside the server, RAM 0, RAM1, CPU0, CPU1, HDD, Ethernet and SP. In addition, there is a set of tachometer sensors to measure and control the servers fan speeds and a set of voltage sensors. The sensor readings are also compared to the design limit of each component and a signal is sent to the head node server in case of reaching the threshold limit.

3.3 Server level cooling components

Traditionally, the Sun Fire V20z servers were designed to be fully air cooled However, in the current configuration, all heat sinks and CPU fans are replaced by direct contact liquid-cooled technology heads manufactured by CoolIT Systems [113]. This will be explained in detail in the following two sections.

3.3.1 Air cooled Sun Fire Servers and servers fans

The Sun Fire V20z servers were cooled using 4 sets of 2 fans distributed inside the server as shown in figure 3.3. The fans used in these servers have casing dimensions of $40 \times 40 \times 28 \ mm^3$ and the specification is listed in table 3.3. The server fans are the key component in temperature dependent power variation as they have variable speeds driven by the platform cooling demand. Thus, the fans can operate over a range of power draws depending on the

thermal load. The fan speed control (FSC) algorithm is fairly complex in the server power design, but in general, a number of different temperatures are measured in the server platform and used to adjusts the server fans accordingly to ensure all components are at or below their upper thermal threshold limits. The FSC generally attempts to run the fans at the lowest possible speed to minimise fan power consumption and noise. Increasing the platform stress load will generally increase the fan speed to reach a required cooling target.
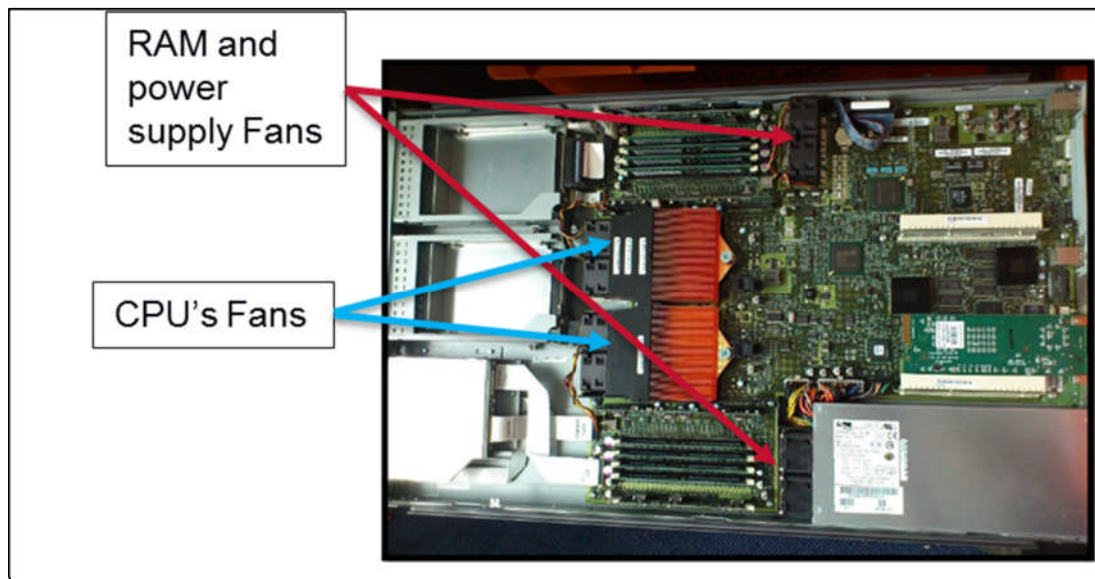
Figure 3.3 The Sun Fire V20z fully air cooled server.

Table 3.3 Specifications of the fans used in Sun Fire V20z servers [114].

Model	Rated voltage	Operating voltage range	Input current	Input power	Speed	Max airflow	Max air pressure	Noise
Part number	VDC	VDC	Amp	Watt	R.P.M.	$\dfrac{m^3}{min}$	mmH_2O	dB-A
FFB0412SHN	12	4.5 to 13.8	0.5	6	13000	0.68	27.3	54.5

The first set of these fans are attached to cool the power supply unit (PSU) and RAM0 which are associated with the master CPU (CPU0). These fans are constantly running with the pre-operation status at a fan speed of 4150 rpm fan speed and increased up to 11600 rpm during the normal operation of the server. Thus, the air passes through RAM0 carrying their heat before passing through the PSU of the server. The second set of fans is used primarily to cool RAM1 which are associated with the slave CPU (CPU1). These fans generate a negative pressure to supply an air flow passing around and through the HDD disk and then through RAM1 to manage their temperature.

The third and fourth set of fans are used to cool the two AMD CPUs of the server. These fans generate an air flow to pass through the fins of the heat sinks which are sitting directly on the top of the CPU as shown schematically in figure 3.4.
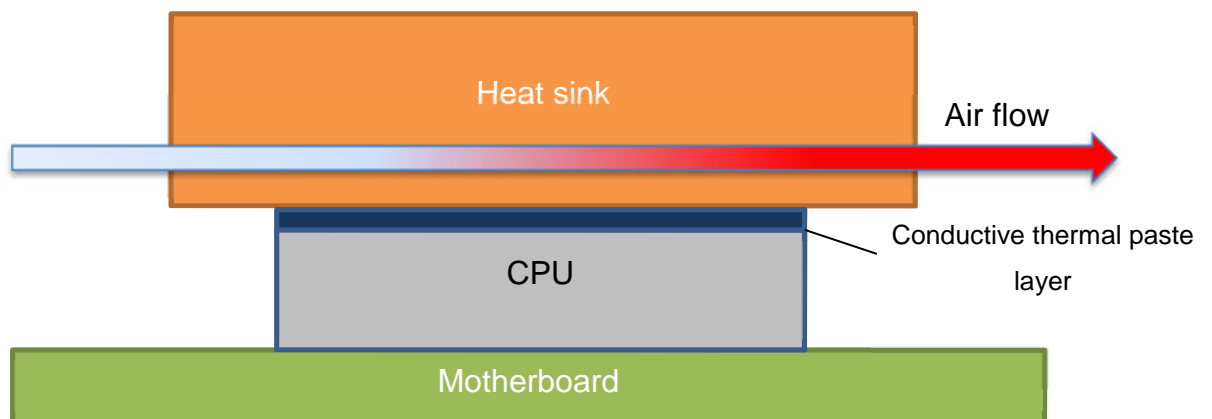
Figure 3.4 sketch of the aircooled CPU.

3.3.2 Direct liquid cooled servers and head pumps

The CoolIT System heads (will be referenced in this thesis as Direct Contact Liquid Cooling (DCLC) heads or servers' modules) were retrofitted to high performance computers (HPC) based Sun Fire V20z air cooled servers. Each CoolIT System head has a small internal pump integrated into it, which can be configured to operate using power from the server or alternatively to centrally pump the coolant around the secondary loop using the pumps in the CHx40 at the top of the rack. With the DCLC heads, microprocessors conduct

heat from the die through thin thermal conductive interface layers to the liquid cooled-cold plates as shown in figure 3.5, replacing the air-cooled heat sinks. The liquid cooling loop accommodates the thermal management of the CPUs although heat generating components like RAM, PSU, hard disk drive (HDD) and other auxiliary components on the printed circuit board are air cooled. Al-Anii's [115] work showed that with the retrofitted CoolIT system heads, there is about 75% of the server heat is transferred into the liquid while the remaining heat is transferred by air. For this, the two sets of fans used to cool the RAM and the PSU are kept in the current configuration and thus the servers are hybrid liquid-air cooled as shown in figure 3.6. Also, the yellow arrow marks shown in figure 3.6 indicate the direction of air flow within the server that is responsible for cooling the auxiliary components.

The DCLC system, shown schematically in figure 3.7, consists mainly of micro fin cold plates, two small pumps, head casing, magnetic mounting kits, flexible corrugated tubes, quick valves connectors (Stäubli type) [113]. These components will be explained in the following points:
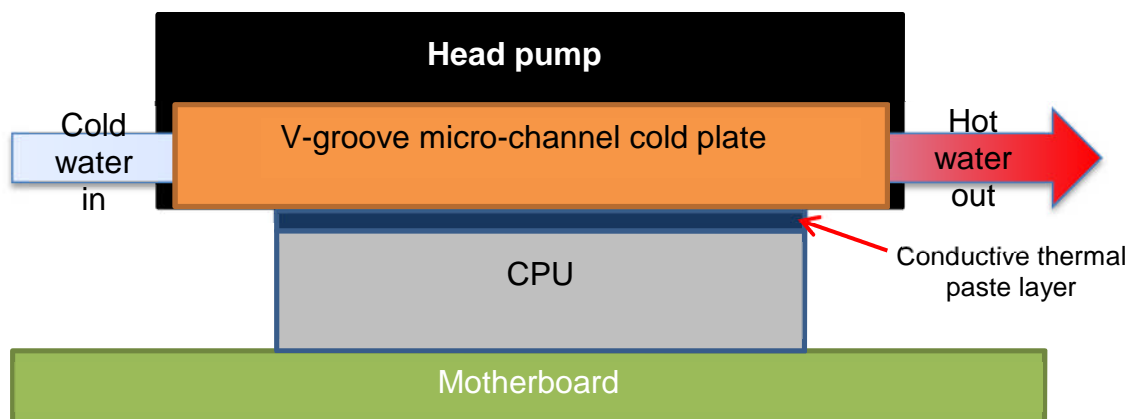
Figure 3.5 Sketch of the DCLC cooled CPU.
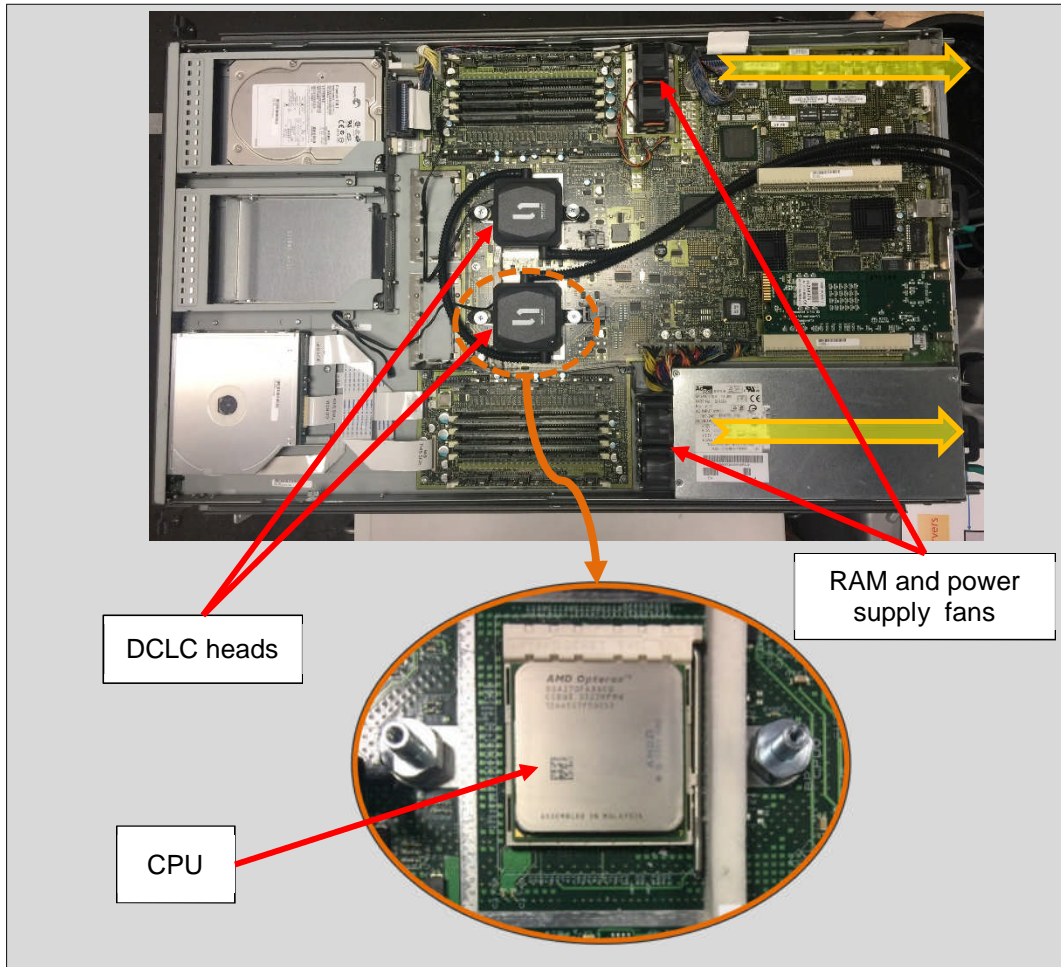
Figure 3.6 Retrofitted Sun Fire V20z server with DCLC units.
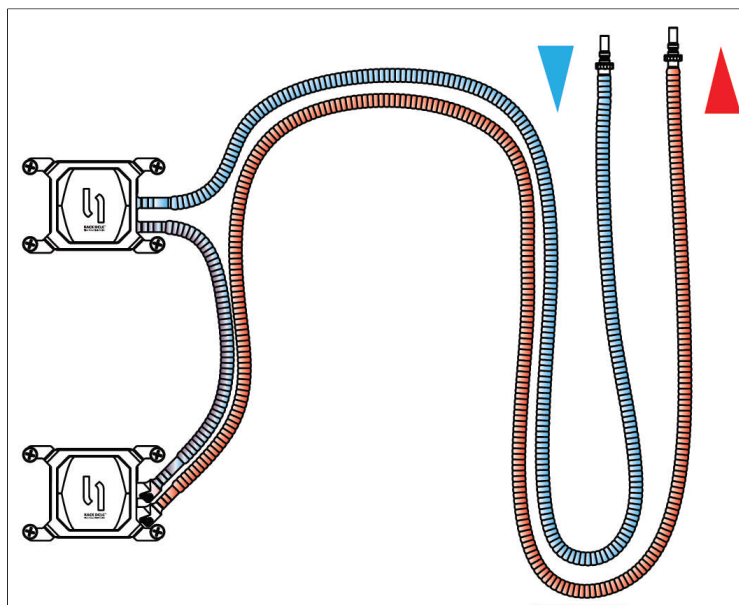
Figure 3.7 Schematic of the DCLC head pumps [113].

65

1. Head pumps and the casing

The DCLC head pumps use brushless electronically commutated motors which uses an integrated inverter power supply to convert the direct currents (DC) into alternative current (AC). The components of the pumps are made of semi-crystalline engineering thermoplastic and Poly Butylene Terephthalate. The case is made of Acrylonitrile Butadiene Styrene where the pump assembly is encapsulated inside. The material selection of the pumps is based on the low flammability rating, high toughness, easy formation, and structural reliability [78]. The nominal rated power consumption and flow rate of the DCLC pump are 1.44W and 0.57 l/min, respectively. The acoustic noise of the pump is relatively small which is about 23dbA at the nominal pump speed of 2200 rpm. Figure 3.8 shows the DCLC pumps and the casing, and figure 3.9 shows the nominal operational curve of the pump.
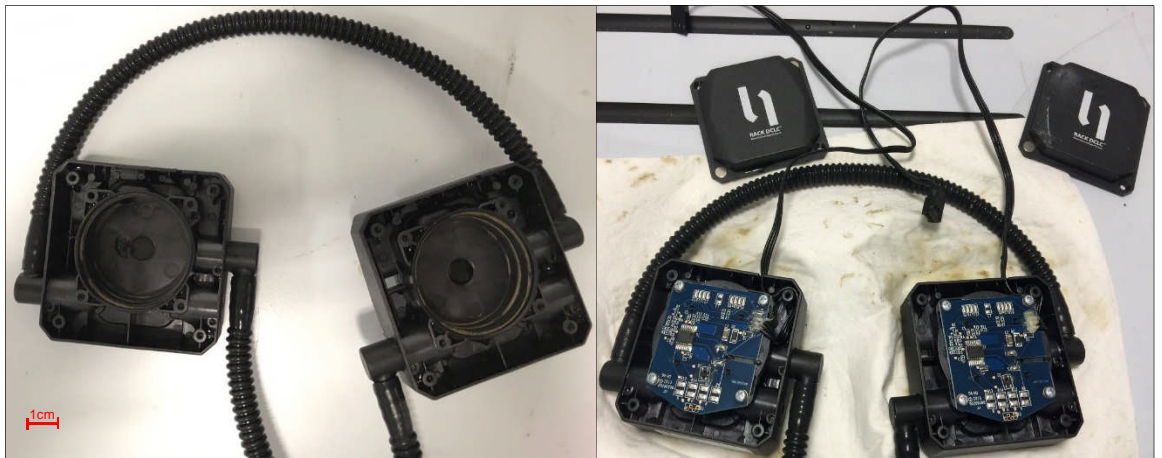
Figure 3.8 DCLC head pumps and the casing.
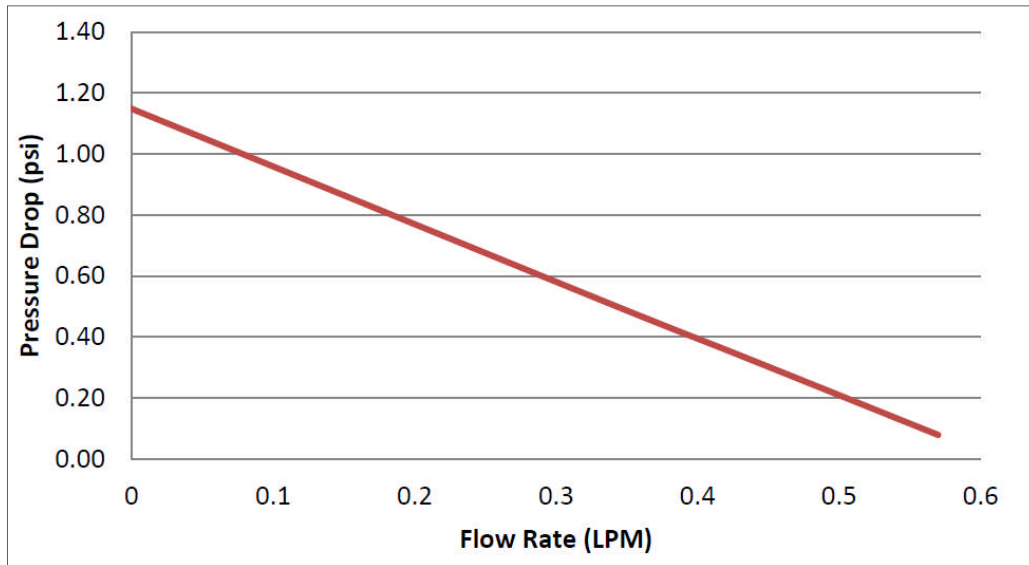
Figure 3.9 Nominal operational curve of the DCLC head pumps [113].

2. V-groove micro-channel cold plate

The cold plate is patented technology of CoolIT Systems and consists of a solid copper component that maximise the performance for high thermal density applications. The cold plate sits on the top of the CPU and is separated by a thin layer of thermal conductive paste. Figure 3.10 (a and b) shows the cold plate heat exchanger and the bottom side of the casing. A rubber gasket is used to direct the coolant to the warmest region of the processor first which is at the core of the CPU as shown in figure 3.10a.

Figure 3.10 (a) DCLC casing from the CPU side showing the cold plate and the attached rubber gasket flow guide and (b) is the V-groove micro-channel cold plate.

3. FEP tubing

The connecting tubes are made of Fluorinated Ethylene Propylene (FEP). FEP is ultra-high reliable, compact, durable, non-flammable, and chemically stable. The FEP tubes connect the CPU head pumps with the intake and discharge manifolds and they have both flexible corrugated and straight

sections. The thickness of the tubes is 2 mm and the external diameter is 8 mm which have strong and durable properties.

4. Quick connector Stäubli valves

Stäubli quick connector valves are used to connect the tubes to the manifold or the heat rejection system as sown in figure 3.11. These valves are especially constructed of durable treated surface metal alloys which gives an extremely robust mechanism with high reliability [113].

Figure 3.11 stäubli valves configuration to connect the heat pumps to the manifold [113].

5. Server loop coolant

The coolant in the servers' loop (the secondary loop of the CHx40) is OAT 25 which is a mixture of 74% deionised water, 25% propylene glycol, and 1% inhibitors. The 1% inhibitors is mainly a mixture of biocide, azule, and anti-corrosion inhibitors. The mixture gives a freezing point of -6℃ and specific heat capacity of 4.075 kJ/kgK.

3.4 Direct liquid cooled rack level design

The rack design includes one rack frame, cloud computing servers, network switch, power shelf, rack power supply unit (PSU), coolant manifold, CHx40 liquid/liquid HE, and a rear door air/liquid HE as shown in figure 3.12. Each item will be discussed in detail in the following sections.
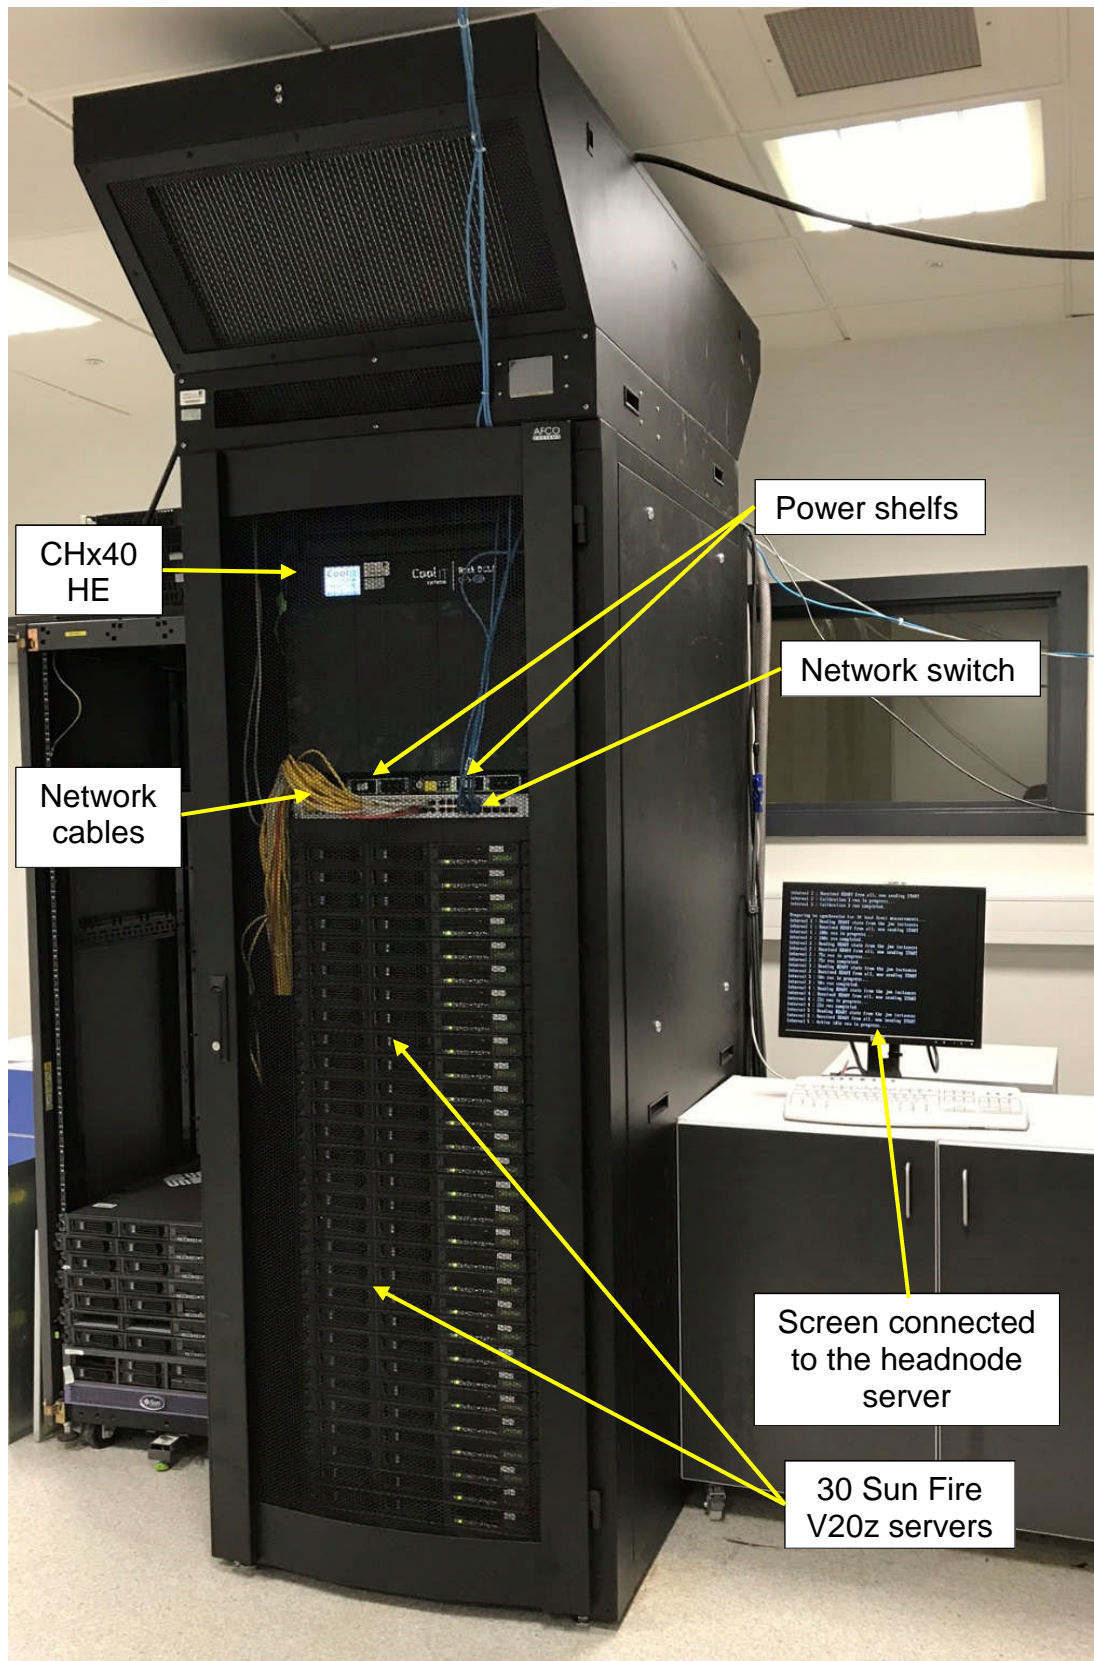
Figure 3.12 Cloud computing rack level.

3.4.1 Cloud computing servers

Thirty Sun Fire V20z servers are retrofitted into a single rack, as shown in figure 3.12. The rack is placed in the lab in the first floor of the Energy building of the University of Leeds. The OS software (version 13.10 of Ubuntu) was installed in the servers by Al-Anii [115].

The top server in the rack was assigned to be the master server to control the remaining servers of the rack and it is called the head node server. All the thirty servers are served up network addresses through Dynamic Host Configuration Protocol (DHCP) over a network switch as shown in figure 3.12. The head node server runs a script which takes Media Access Control (MAC) addresses that gives a unique identifier address to all of the servers and generates names and IP addresses. Thus, the head node server is called eng01, and the second server from the top is called eng02 and so on for all the thirty servers. The IPMI is accessed for all the servers via daisy chain data cables connected in series to the management ports of the servers and to the top of rack switch. The servers are all connected to the network switch for local hosting networking and internet.

3.4.2 Network switch

The network switch from 3C0M [116] includes 48 10BASE-T/100BASE-TX/1000BASE-T network ports and 4 SEP ports for management access and various system status indicators. The switch has a standard AC power socket (100-240 VAC) in the rear panel and is directly connected to the PSU of the rack. It is mounted on the top of the servers using two brackets to keep the switch in a horizontal position. All the servers are connected to the switch for networking and internet. It also allows to control the rack remotely, for running stress testing, downloading the data and monitoring the operation.

3.4.3 Power shelf and rack PSU

Two power shelf switches are used to power the servers as shown in figure 3.12. These switches are connected to two PSUs. So the servers, network switch, and the CHx40 heat exchanger are powered by two PSUs. The two

power shelf switches are located over the network switch to allow easy access from the front of the rack.

3.4.4 Distribution and collection Manifolds

After all the servers have been placed in the rack, the DCLC cooling loops of all the servers are connected using the quick connections manufactured by Stäubli at the rear of the rack to two plenums that work as the manifolds. The cooling loops are joined together through two passages: one for the cold side (coloured blue) which represents the inlet to the servers and the other one for the hot side (coloured red) which represents the rejected water from the servers. These manifolds are connected to the secondary loop of the CHx40 HE. The manifolds are six feet long each with square cross section (25×25mm internal dimensions) made of stainless steel. The manifolds employ dry-break quick connection technology and can accommodate 42 servers modules connected in parallel as shown in figure 3.13.
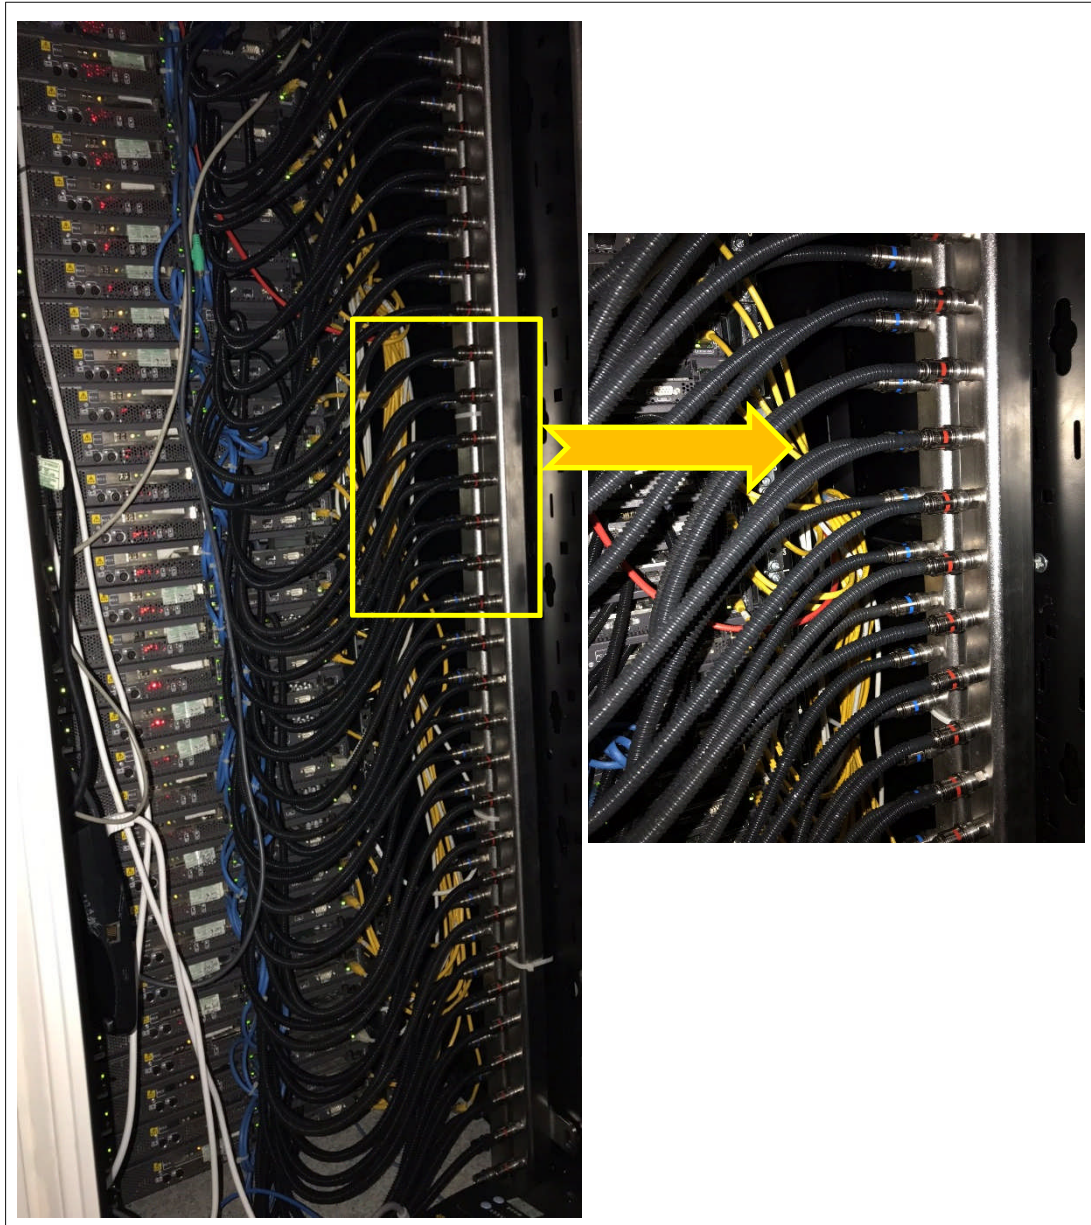
Figure 3.13 Back of the rack showing the manifolds and DCLC stäubli valves connections.

3.4.5 CHx40 heat exchanger

The CHx40 heat exchanger module is designed by CoolIT Systems [113] with a height of 2U (88.9mm), for distributing the coolant through a series of DCLC cold plates situated inside servers within a rack. The thermal load is exchanged in the CHx40 module from the secondary fluid to the primary fluid using a liquid to liquid heat exchanger. The CHx40 module is installed in the

very top point of the rack and connected with the manifolds in the rack and the AHU using the stäubli quick connectors SPT12 and SPT10, respectively.

Essentially, there are two cooling liquids which operate the cooling cycle in the CHx40; the liquid that flows through the DCLC cold plates into the server which is called as the secondary liquid and the liquid with which secondary coolant exhausts its heat across the heat exchanger is called as primary liquid (process water in the AHU). The filtration requirement in the primary loop is 300 microns, for this, a filter is designed and fitted in the loop in the AHU configuration design which will be discussed in details later in this chapter. The CHx40 consists mainly of plate to plate heat exchanger, two pumps, valves, fittings, reservoir, temperature sensors, pressure sensors on both loops, flow sensor on the secondary loop, and humidity sensor, as shown in figure 3.14. The two pumps are connected in series on the secondary loop to increase the pumping head and for n+1 redundancy.
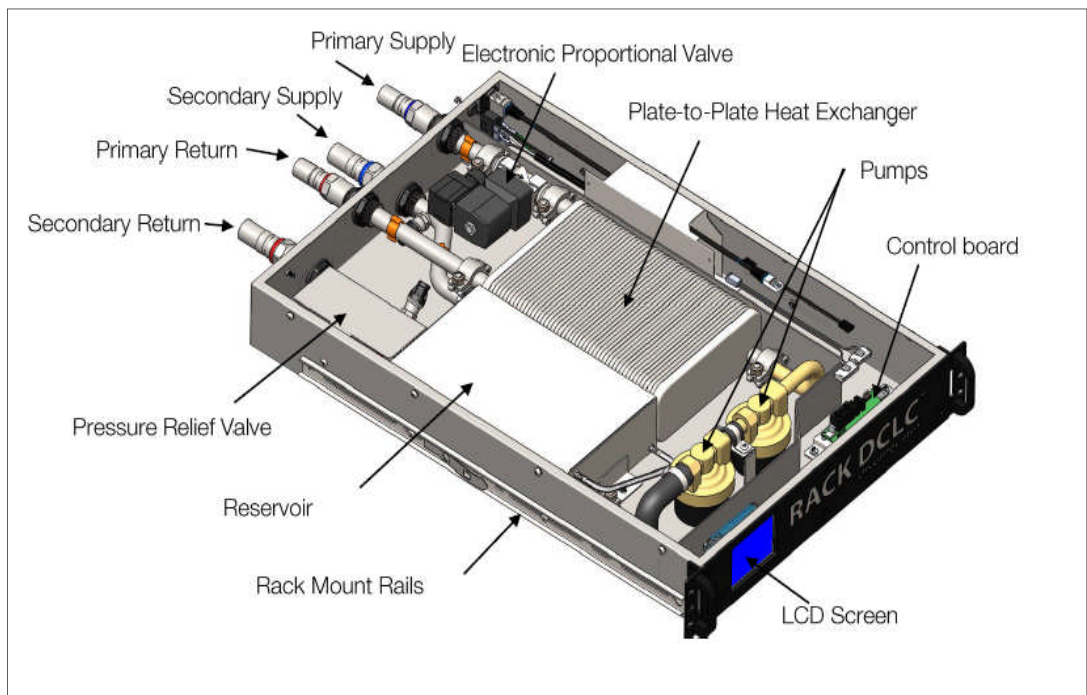
Figure 3.14 CHx40 liquid to liquid heat exchanger [113].

A proportional valve is equipped on the primary water loop circuit to control the process water flow rate through the CHx40 module. The proportional valve setting can be controlled from the LCD screen on the CHx40 and it can be

adjusted to operate automatically to keep the process water temperature higher than the dew point temperature of the IT room to reduce the risk of condensation inside the servers.

The designed cooling capacity of the CHx40 at 30℃ inlet temperature and 35l/min flow rate of the primary coolant is 40kW. The cooling capacity of the CHx40 as a function of the inlet primary water flow rate for different inlet temperatures is shown in figure 3.15.
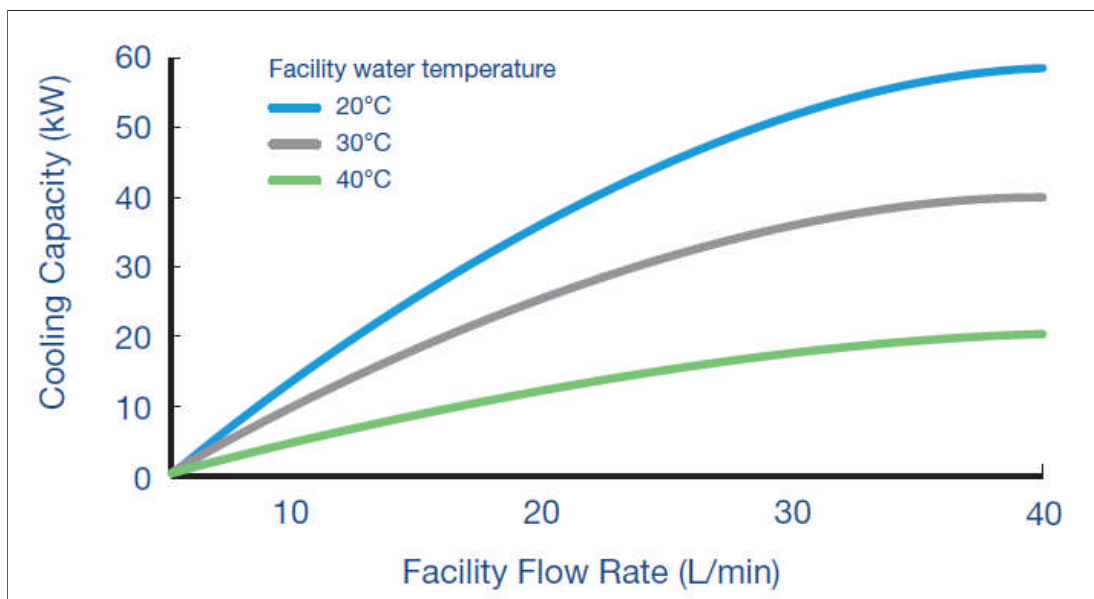
Figure 3.15 Operation curves of the CHx40 for various process water flow rate and inlet temperature [78].

The CHx40 operates in the voltage range of 100 – 240 V AC and the maximum power consumption is 160 W. It is also powered from one of the rack PSUs. The CHx40 is also connected to the head node server `eng01` in the rack through the network switch for data logging. The control system module is a Linux barebone based computer running Angstrom distribution. Hence, the control system can be connected to the head node via an SSH client with the assigned IP address and data can be read that appears on the status screen on the LCD of the CHx40.

3.5 Rear door heat exchanger design

To collect the heat generated by RAM and the PSUs of the servers, a passive rear door heat exchanger (RDHx) was installed. This device consists of a wavy fin and tube heat exchanger which is mounted in the rear door of the rack. Chilled water supplied from the AHU flowing through the heat exchanger tubes is used to cool the exhaust hot air which is expelled from the rear of the rack. The RDHx offers many features including eliminating local hot spots occurring in the data centre, low total cost of ownership, high reliability, ease of installation.

The original rear door of the rack was replaced by a copper wavy fin and tube heat exchanger ($16 \times 61 \times 200 \ cm^3$) which is fitted into the rear section of the rack as shown in figure 3.16 and then customised to fit the rack dimensions to guide all the air generated by the server fans to pass through the RDHx fins.

The RDHx design was designed to function without any added air moving devices or controlling devices and relying on server fans only (passive way of cooling). Hence, the design of the tube size, fin shape and spacing, had been chosen to minimise the air impedance through the heat exchanger. The inlet process chilled water to the RDHx was designed to be maintained above the dew point temperature at all times to eliminate the chance of condensation and the need for a complicated design to drain the condensed water. This was achieved by connecting the RDHx in series with the primary loop of the CHx40 which contains a solenoid valve to shut the flow when the flow temperature drops below the room ambient temperature.

The process chilled water is delivered to the RDHx using flexible (25.4mm inner diameter) braided type hoses. The hoses are brought close to the bottom hinge axis of the RDHx and designed to rotate around each other to allow the RDHx to be opened and closed with minimal movement to the braided hoses to allow normal access to the rack interior.
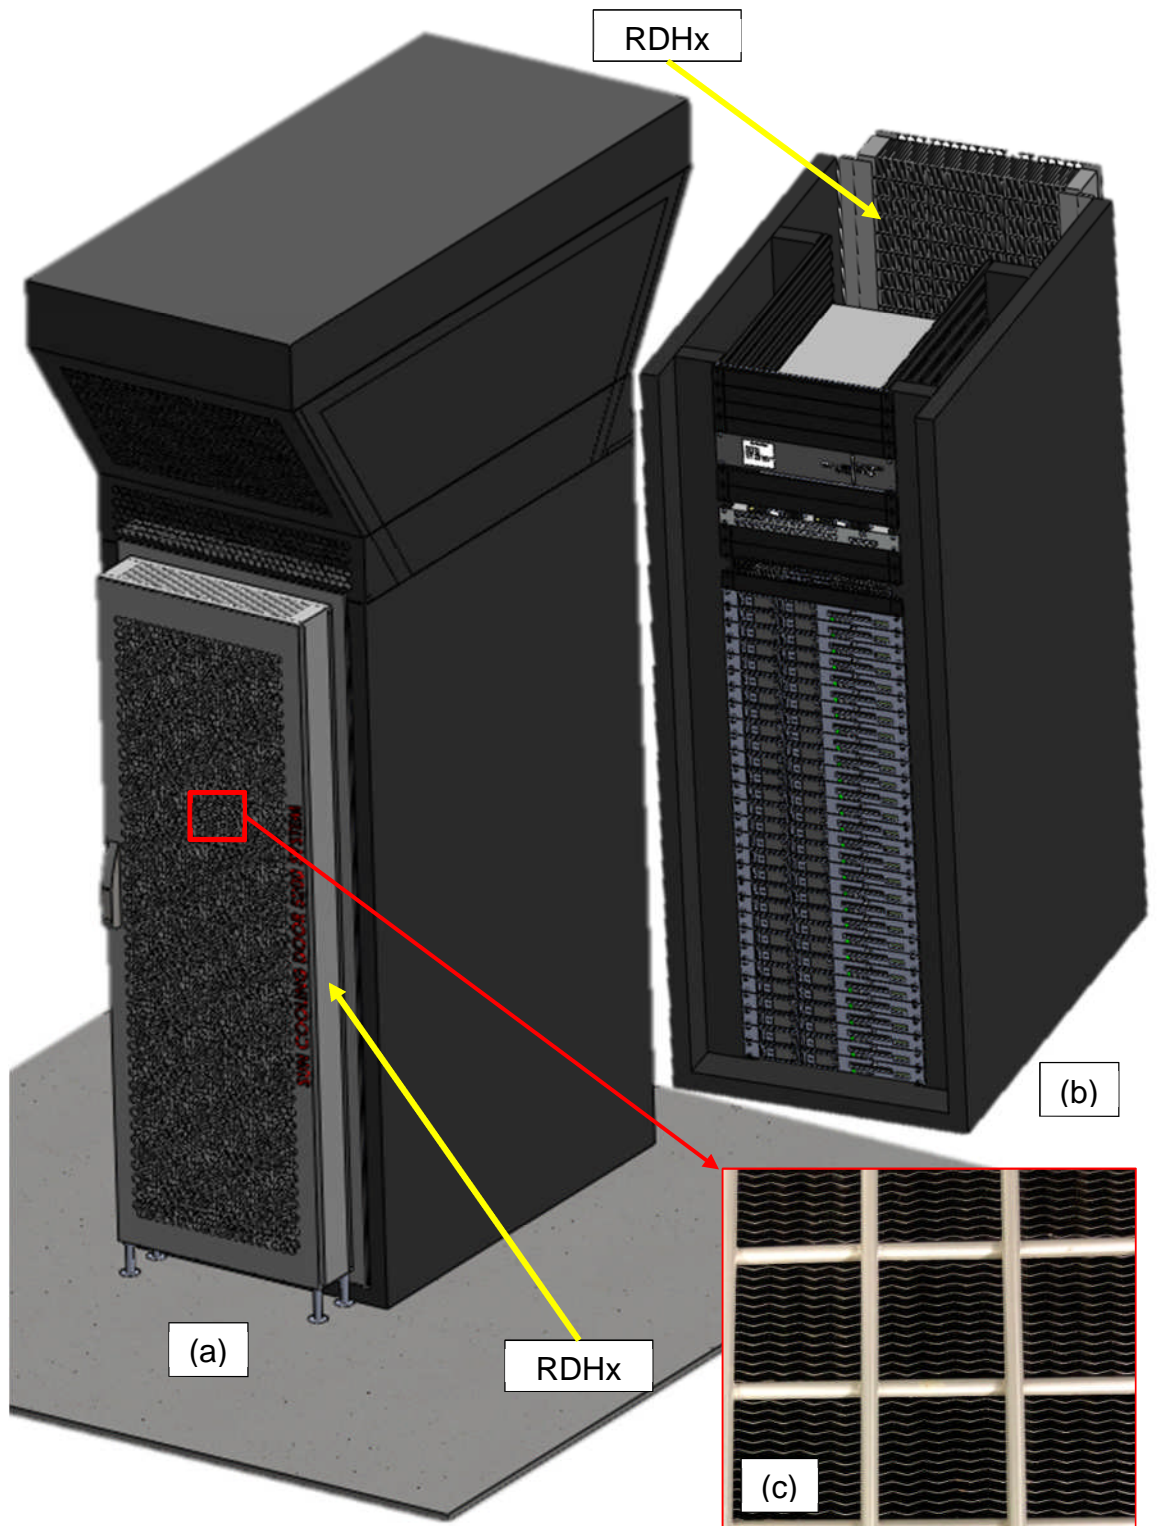
Figure 3.16 Rear door heat exchanger design (RDHx); (a) shows the back view of the rack where the RDHx is attached to the rack, (b) shows a section view of the rack and (c) shows the wavy fins of the RDHx.

3.6 Air Handling Unit (AHU) design

The AHU was designed as part of the present work and manufactured in collaboration with Airedale International Air Conditioning Ltd [117]. The main purpose of the design is to investigate the chiller-less and innovative design of AHU that utilise evaporative cooling in direct liquid cooled data centres. The designed AHU is placed in the plant room of the Energy building of the University of Leeds. The AHU has exhaust heat ducted into a large ducting system for the Energy building. The AHU operation is based on the ambient air conditions in the plant room, which can be in excess of 25℃ in the summer and sub-zero in the winter. Figure 3.17 shows a full picture of the AHU which consists mainly of three loops:
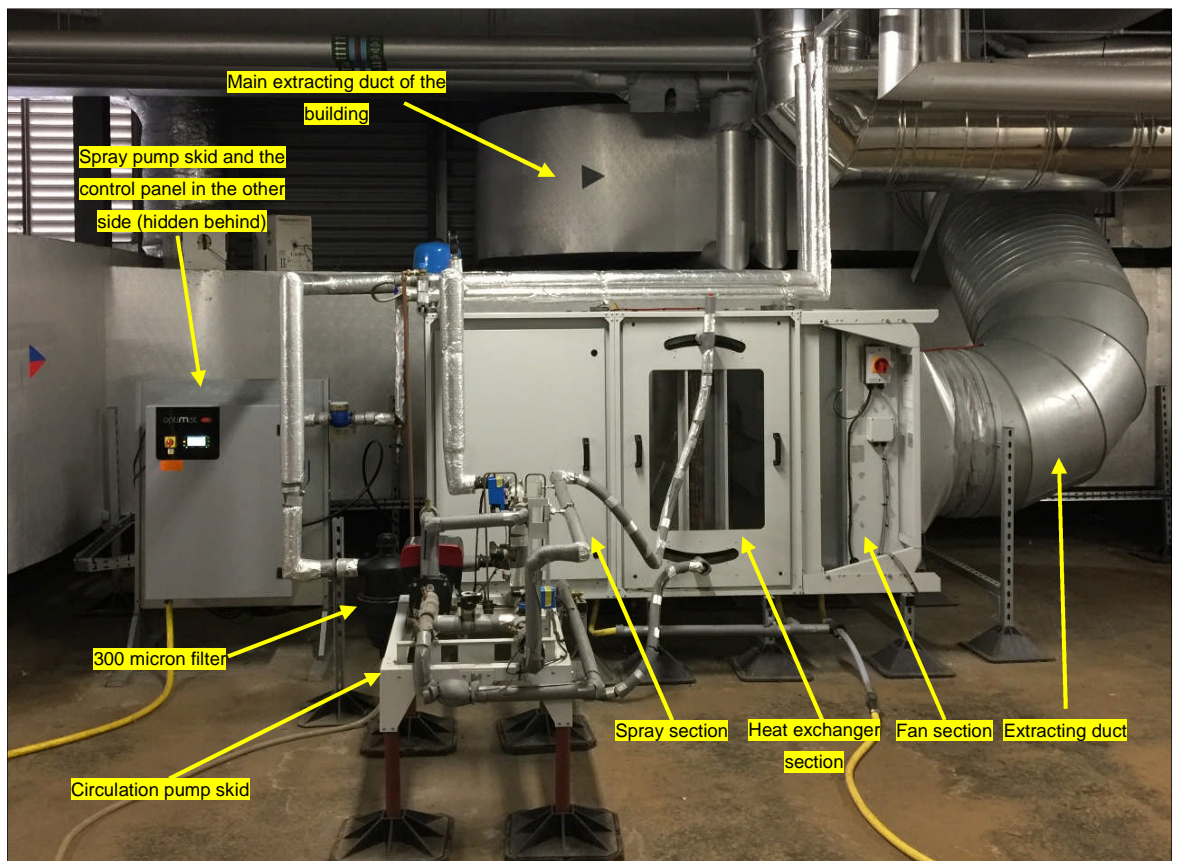
Figure 3.17 The designed bespoke air handling unit.

1. Processing water loop: the water coming from the CHx40 enters in hot condition to the AHU HE where it rejects the heat before being pumped back to the CHx40. A pump, filter, bypasses loops, pressure vessel, and different sensors are all included in the loop.

2. Airside loop: The rejected heat is carried by the passing air through ventilation ducting to the external environment. The AHU was designed as an open circuit wind tunnel of 1103 X 1197 mm^2. This loop consists of Inlet air section, heat exchanger section, visualization section (perspex panels are used to view into the HE section), axial fan, and different sensors.

3. Spray water loop: To enhance the cooling capability, a spray system has been bolted onto the AHU HE giving more than 15kW cooling capability. This loop includes spray pump skid, filter, control system, valves, different sensors and nozzles.

The system is remote monitored (over ethernet) and controlled using a pCOweb card (essentially a linux based ARM system).

The design characteristics of all the components, the full experimental setup, instrumentation and control will be related in details in the following sections.

3.6.1 Heat exchanger design and selection

Compact round tube heat exchangers are typically used as condensers for many industrial and commercial applications. However, the performance of a dry cooled heat exchanger (HE) degrades considerably with increase in ambient air temperature and in some applications the typical HEs are not capable to carry the required cooling load. To overcome this issue, it is necessary to explore the methods to improve the cooling performance and select critically the main parameters which are involved in the HE operation. The enhancement can be achieved by utilising extended surfaces, geometry modifications, and surface wettability and in wet conditions by utilising water evaporation in the air side of the HE.

A number of experimental studies have published performance data of dry cooler HE and also found correlations for predicting air side heat transfer and pressure drop performance in terms of Colburn j and friction f factors as a

function of HE geometry and Reynolds number [118-127]. In comparison to round tube HE performance data in dry conditions, published experimental work in wet conditions in the open literature are limited in scope [128-131]. Thus, researching the impact of fin geometry on the dry and spray evaporative HE from the literature will be informative for developing design guidelines and to help realize the full potential of using spray droplet with the fin surface in HE industry, specifically in data centres.

The literature has revealed that there are many different parameters that affecting the performance of round tube HEs. However, the main parameters are summarised in table 3.4 to help in selecting the proper heat exchanger for the AHU.

Based on table 3.4, the heat exchanger is selected to be round tube plain fin (RTPF) type. It has tubes running end to end, through aluminium fins which increase the surface area. The fins are epoxy coated to protect them from corrosion. The copper tube diameter is 9.52mm (outer diameter) and the thickness is 0.28mm. The fin pitch is 1.8 mm and the thickness is 0.11mm.

Table 3.4 Heat exchanger selection parameters

Parameter	Reference	Critical finding	Present study HE design selection
Fin Shape Plain, Herringbone wavy, Sinusoidal wavy, and Louvre	[121, 132, 133]	In dry conditions: • $h_{wavy} > h_{plain}$ [55%-70%] $f_{wavy} > f_{plain}$ [66%-140%] • $h_{louvre} > h_{plain}$ [99%-180%] $f_{louvre} > f_{plain}$ In wet (spray & deluge) conditions: Very limited number of studies regarding this and even though the available literature data refer that there is no much heat transfer improvement over the pressure drop penalty. (h is the heat transfer coefficient)	Improving the shape of the fins in fin and tube HEs generally show better heat transfer characteristics. However, this improvement is usually accompanied with higher manufacturing costs and pressure drops (friction losses) which in turn increases the power consumption in the fan. On top of that, in the wet case where the spray cooling is utilised, the friction losses will dramatically increase with using complicated fins such as the louver and wavy fins which will increase the power required by the fan to achieve the targeted air flow rate. In conclusion, plain find round tube HE will be chosen to be fitted in the AHU to achieve lower cost of ownership and operational cost.
Fin Pitch & Number of Rows	[120-122, 125]	In dry conditions, the effect of fin pitch is negligible for the number of rows 1 & 2. However, decreasing fin pitch has a notable increase in the pressure drop and heat transfer performance when the number of rows more than 3. ($h_{Fp=1.7mm} > h_{Fp=1.2mm}$ [42 %]) In condensation wet conditions, changing fin pitch has a high effect on both heat transfer and pressure drop characteristics even for small number of rows (2&3).	According to the literature, choosing the correct number of rows in HEs is very crucial especially in the wet case, higher number of rows is higher friction losses and in turn higher power consumption of the fan. For this, the number of rows of the HE exchanger is chosen to be two for better performance and lower power consumption. The idea of decreasing the fin pitch is to obtain higher surface area of HT and thus small HE size. However, this is usually accompanied with higher pressure drop especially in the wet cases where the effect is very high as shown in the literature. For this, the fin pitch is chosen to be 1.8 mm and surface area of 1105×870 mm^2.
Effect of Inclination	[129, 131]	In dry conditions, the effect of inclination is negligible. In wet deluge cooling, $h_{\theta=21} > h_{\theta=0}$ [8.3 %] (θ is the HE inclination angle) $f_{\theta=21} > f_{\theta=0}$ [2.2%] Angle 20° is found to be the most practical which is needed in applications for space and performance constrains.	The AHU design of the present study allows inclining the HE between 0° and 20°. However, this aspect will not be used in the present configuration of the study and will be suggested for future work. Hence, the HE will be kept vertical.
Effect of Coating	[126, 128, 131, 134, 135]	In dry conditions, no change in h and f with coating the HE. In wet condensation conditions, • $h_{uncoated} > h_{coated,hydrophilic}$ [10%] • $f_{uncoated} > f_{coated,hydrophilic}$ [40%] Types hydrophilic treatments: • Sodium hydroxide and potassium persulphate • Plasma In evaporation conditions, For deluge cooling • $h_{uncoated} < h_{coated,hydrophilic}$ [40%] • $f_{uncoated} < f_{coated,hydrophilic}$ [15%] For spray cooling • $h_{uncoated} < h_{coated,hydrophilic}$ [14%] • $f_{uncoated} \approx f_{coated,hydrophilic}$	As the effect of the type of coating is not fully established in fin and tube HE and it is very costly in our experiment due to the practical large scale system we use which requires coating large scale HEs. However, the HE used in the present study will not be coated and will be used with the typical epoxy coating. However, a small scale study of the effect of coating is achieved for different type of coating for a single droplet. The study could not mimic the real characteristics of the HE but will show the effect of the type of coating on the droplet evaporation for the same temperature.
Way of utilising evaporative cooling	[131]	A single study was found regarding Herringbone wavy fin in round tube HE under two different types of water distribution: deluge and spray types. It is found that spraying the HE gives better characteristics regarding pressure drop and less water usage.	Based on the results of [131] for Herringbone wavy fin type under adiabatic cooling, the spray technique is chosen rather than deluge in the present design of the AHU.

3.6.2 The EC axial fan

An axial fan from ebm-papst [136] is chosen to be used to cool the dry cooler heat exchanger with air passing through it. The design of the axial fan blades and its rotor motor allow compact dimensions, high air performance at minimal noise. The fan blades use a support structure made of aluminium alloy of high robust and corrosion resistance. The blades are coated with special fibre reinforced plastics. These properties of the blades help to use the EC axial fan in the present design of the AHU where an adiabatic cooling is utilised.

The type of the fan used in the AHU is W3G630-GN37-23 and the nominal voltage and power consumption is 230V and 730W respectively at 960rpm which is equivalent to 8100 m^3/h. The external frame dimensions of the fan are 805 × 750 mm^2 which fits the frame of the AHU. The fan operation curve is shown in figure 3.18.
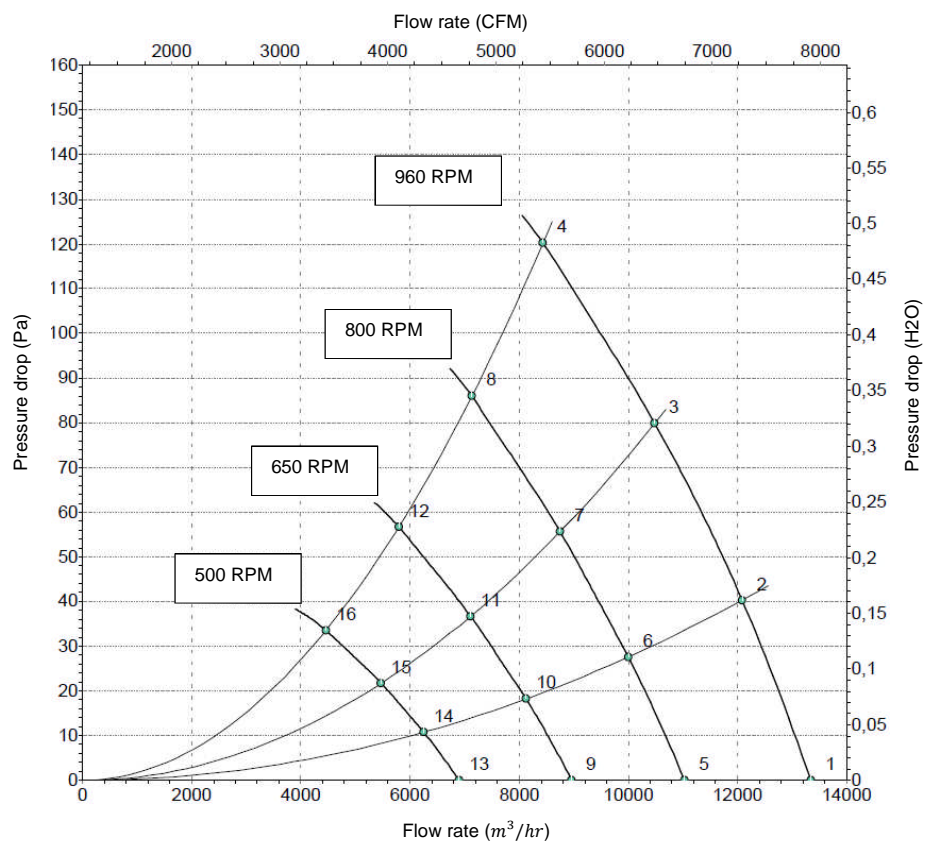
Figure 3.18 Fan operation curve of the AHU (the numbers representing different measurement points to obtain the operation curve) [137].

3.6.3 Spray system

To increase the AHU performance and maintain a reliable operation of the unit at challenging environmental and data centre operational conditions, an adiabatic spray evaporative cooling is employed. For this, an evaporative spray cooling system called optiMist from CAREL [138] is installed in the AHU to cool the air before passing through the heat exchanger or to spray the heat exchanger directly. The spray system atomises water into fine droplets which then evaporate in the air or on the heat exchanger surfaces, removing heat from the heat exchanger. The spraying system is activated under certain conditions of weather and data centre operational load and the spray water flow rate is measured to define the performance.

The optiMist system uses a vane pump to pressurise the water, subsequently atomising it through special nozzles. The sophisticated control system of the optimist combines the job of the inverter, which regulates the pump speed and consequently the spray flow rate, with two solenoid valves, ensuring the system always works at the optimum pressure to atomise the water, across a wide range of spray flow rates. Figure 3.19 shows the pumping station and the electronic controller of the optiMist. The pumping station pressurises the spray water to 4-15 bar and manages the spray flow rate for maximum precision, minimum water consumption and minimum energy use. As detailed in figure 3.19, the pump station contains the hydraulic components and the electrical components both in one cabinet. The hydraulic section contains a water treatment system, which is a reverse osmosis demineraliser (to supply low mineral salt water), pressure transducers, manometers, solenoid valves, and pump. The electrical section, contains the power terminals, pCOWeb controller, and the inverter which is the main components to regulate the pressure and the flow rate.
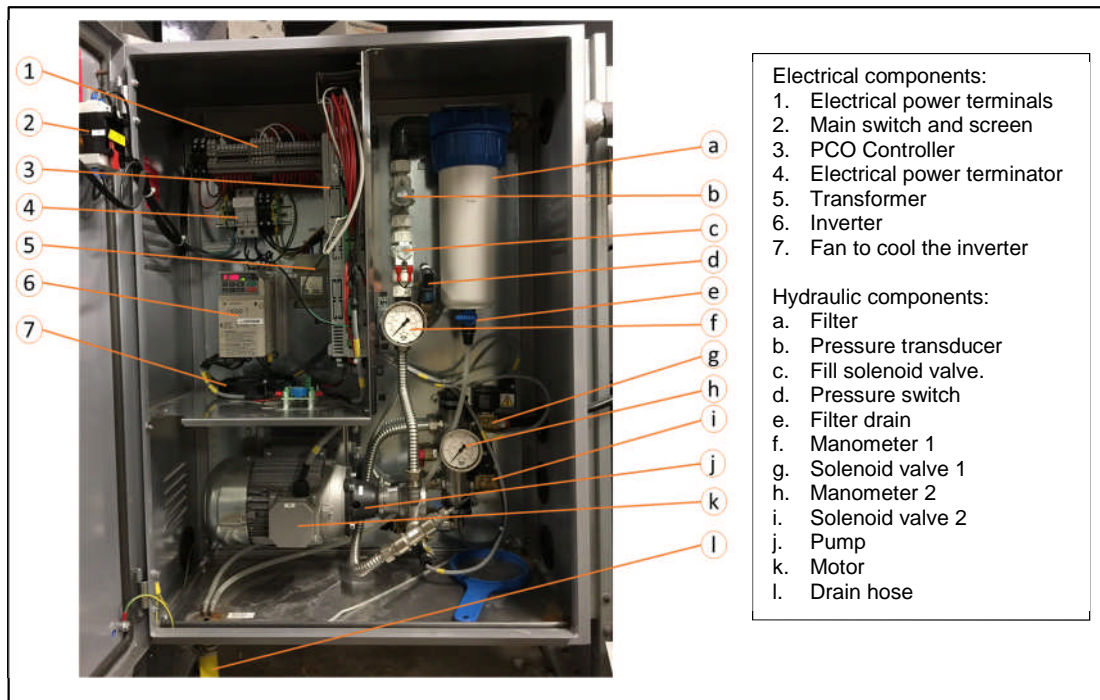
Figure 3.19 Spray pump station.

The nozzles spray system is installed upstream of the heat exchanger, which is formed from six spray nozzles in a rack formation as shown in figure 3.20. Each nozzle has circular orifices and spin inserts to aid with atomisation and formation of a 60° even cone of spray. The nozzles atomise the water to droplet diameter in the region of 15-30 microns. An automatic drain valve was installed on the nozzles rack in order to empty the droplet collection tray when the spray system is not in use.

A rail was specifically designed in the AHU frame to carry the nozzles rack so that the position of the nozzles rack system can be adjusted relative to the heat exchanger position. This gives the flexibility of the spray adiabatic cooling scenarios which involves either spraying the passing air to drop the temperature down before passing through the heat exchanger or spraying the heat exchanger directly. In both cases, the effect of the cooling is due to the evaporation of the droplets of water: the change in state from liquid to vapour occurs by subtracting the energy from the air or the heat exchanger surface that (depends on the spraying scenario), as a consequence, is cooled.
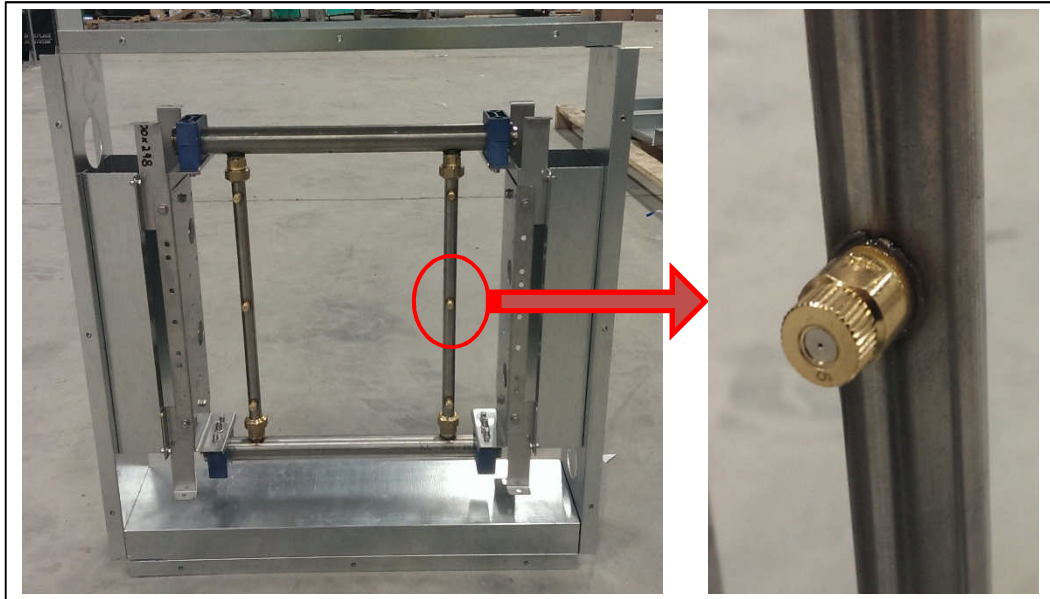
Figure 3.20 spray nozzles in the rack configuration.

3.6.4 Process water recirculation pump

A Grundfos (CME 1-2 A-R-I-E-AQQE) [139] inverter pump was selected to recirculate the process water with the CHx40. The selected pump has to be able to overcome the pressure losses and pump the coolant between the AHU heat exchanger in the plant room (third floor) and the lab where the IT equipment is installed (first floor) which required horizontal pipe lines of $50m$ and $8m$ of pipe line risers. The Grundfos pump is a reliable multistage centrifugal pump with axial suction port and radial discharge port and the pump materials in contact with the coolant are stainless steel which gives a high corrosion resistance.

The pump includes a frequency converter and PI controller in the motor terminal box to enable continuously variable control of the motor speed, which enables adaptation for different required flow rates. The pump is able to work in the data centre temperature range and the nominal calculated flow rate is 18.3 l/min which results in a 21.62m head. Figure 3.21 shows the technical operational curve of the pump.

The pump is installed over a steel frame as shown in figure 3.22 and connected in the output side of the AHU heat exchanger which has the lowest coolant temperature of the loop for longer pump life time. Two bypass loops

were designed to be fitted before and after the pumps, which are used to assist in regulating the temperature supplied to the IT. These bypass loops are remotely controlled and will be discussed later in the section on control design. A filter of 300 microns is installed in the loop after the pump to keep the process water clean and within the required specification of the water loop especially for the CHx40.
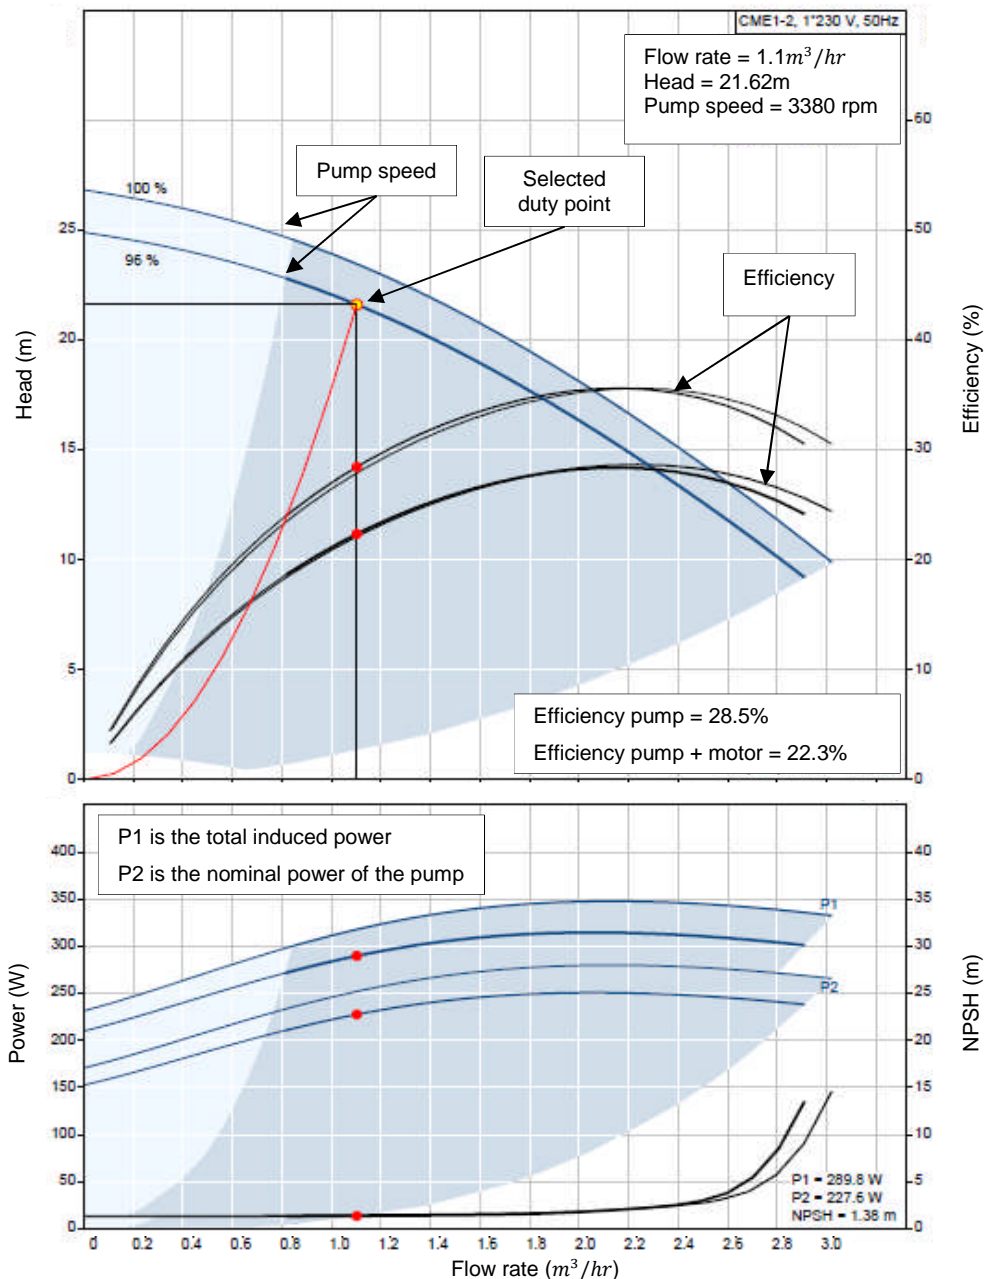
Figure 3.21 Operation curve of the recirculation pump; the dark blue area represent the recommended duty range for best efficiency [139].

86
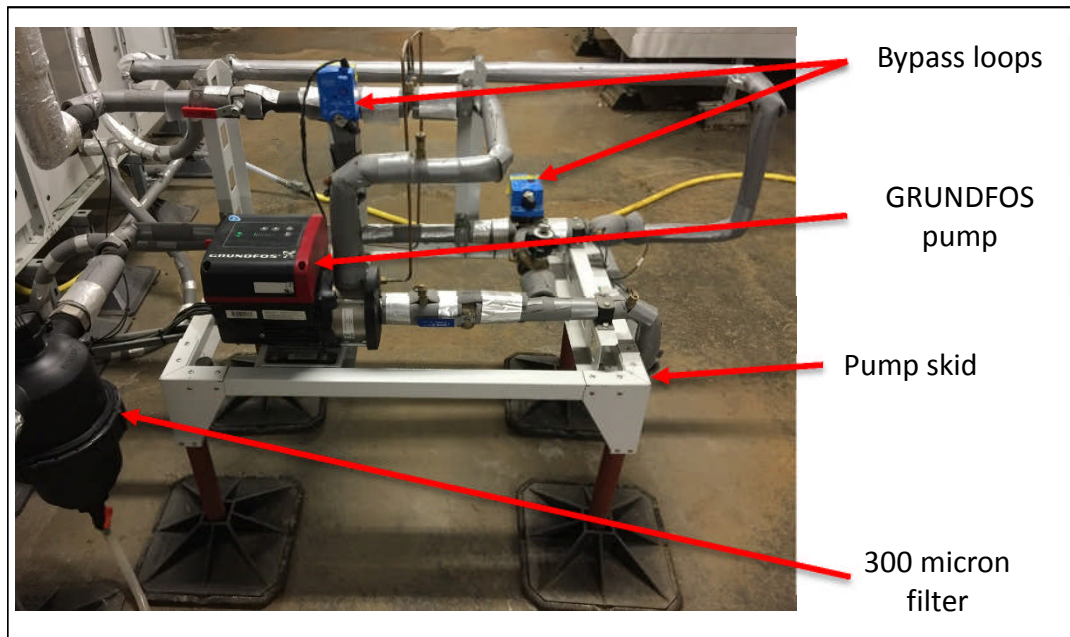
Figure 3.22 Recirculation process water pump.

3.6.5 Control unit and data logging

To perform the study, instrumentation was provided at various locations in the AHU. All sensors of temperature, flow rate, pressure drop, humidity, and power are collected through the pCOweb controller unit, as shown in figure 3.23. The location of the various measurement points will be presented in the next chapter. The AHU can be manually controlled using the customised programme in the pCOweb software: the process water pump (primary loop flow rate), the recirculation valves, the spray system and the AHU fan speed. The software programme also provides a platform for various automated algorithm implementations. The pCOweb is connected to the headnode server (eng01) in the rack via the Ethernet cable which is in turn connected to the university network via a SSH (secure shell) protocol to enable remote control and monitoring.
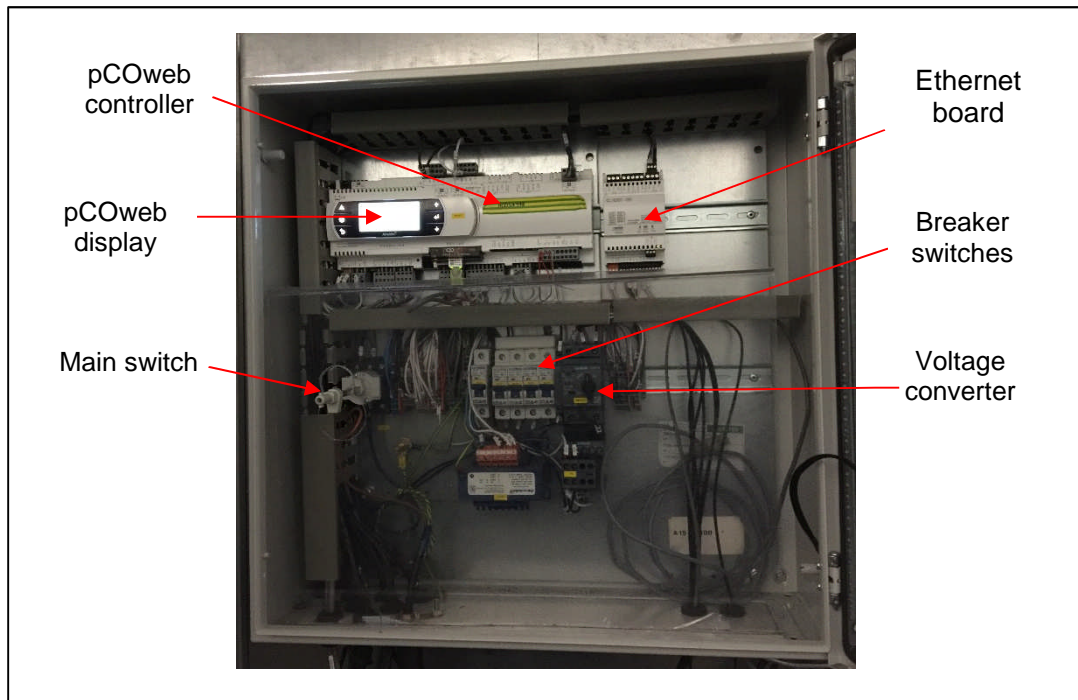
Figure 3.23 Photo of the pCOweb control unit used to control and monitor the AHU.

3.7 Control design of the AHU

3.7.1 Operation information

The AHU is isolated at the main skid control panel. All system components are sub-fused from the control panel. The AHU fan also includes an isolator box on the frame. To electrically isolate the system, the isolator on the main skid control panel can be used.

The system can be disabled temporarily using on/off menu from the control panel screen. This is also possible using the 'Supervisor System Control Status' from the Adjustment screen of the web page remotely. The set points can be adjusted using the pCOweb display. These parameters can also be altered remotely using the adjustments screen on the web page.

3.7.2 Flow control

1. Recirculation pump control

The flow is maintained regardless of cooling demand. The pump is controlled remotely by the pCOweb. The default set point of the flow rate is 0.3 l/s and

can be adjusted. The process variable of the flow rate (l/s), is derived from a differential pressure transducer (kPa), and can be calculated from equation 3.1:

$$Q = \frac{kV \sqrt{\Delta P}}{36} \qquad (3.1)$$

Where: Q is the flow rate (l/s)

kV is the flow coefficient (2.15 m^3/hr at 1 bar differential)

and ΔP is the head loss (kPa)

The flow is PI (Proportional Integrated) controlled at the pCOweb, using 0-10v as an external control signal to the inverter. The control system monitors the status of a normally open contact on the pump contactor. An alarm is generated if the pump status input does not close within 2 seconds of the pump output switching on.

2. IT side Bypass loop control

A proportional valve which bypasses the IT HEs to provide lower flow rates can be manually operated from the display. Opening this bypass allows the AHU HE supply water back onto the AHU HE return and will influence the cooling control. As such an alarm will be generated to warn the operator when the bypass valve is not closed.

3. Permissives

The recirculation pump is permitted to start when the differential pressure transducer is healthy and no inverter alarm exist.

4. Interlocks

The pump is interrupted in the event of:

a. Differential pressure transducer not healthy

b. An inverter alarm is on

c. Loss of pump contactor status

89

3.7.3 Temperature Control

1. Temperature control strategy

The loop is controlled using one of two temperature sensors; at the inlet to the IT HEs or the inlet to the AHU HE. The Process Variable (PV) is selected at the display in the control panel, the temperature is monitored and is PI controlled, maintaining the required set point. It should be noted that when choosing the temperature at the inlet to the AHU HE as the PV of the controller, the AHU HE bypass will no longer be used to reduce the efficiency of the AHU HE and will remain closed all the times.

While the AHU HE provides the bulk of the cooling for the loop, the spray system will assist the AHU HE with adiabatic evaporative cooling toward the upper limits of the PI control band. Also, to avoid the AHU HE over cooling, a proportional bypass valve will begin to open at the lower limits of the PI demand. Figure 3.24 shows the control design based on the chosen PV and demands.
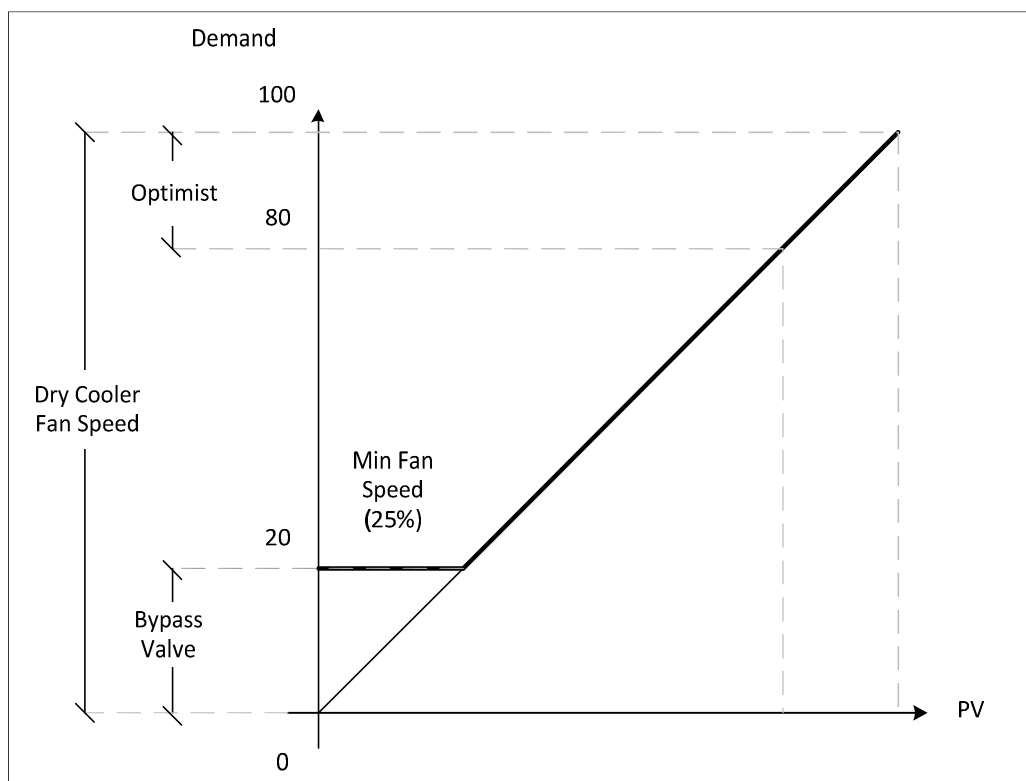
Figure 3.24 Temperature control scheme.

2. AHU fan

The AHU HE operates at all times and when there is no cooling demand there will be a minimum fan speed of 25%. The AHU fan will be forced to its minimum speed when the chosen PV inlet temperature sensor is not healthy.

3. Spray system

The default of the spray system is to operate automatically; however a Hand/Off/Auto switch is present in the controller display in the control panel. In the Hand mode, the spray can be forced to 0-100% of the output. In the Off mode, the modulated output and start commands are disabled. In the Auto mode, the spray will control to the same PV as the AHU fan and the bypass valve. The drain valve is energised in order to empty the droplet collection tray when the spray is not in use to prevent Legionella bacteria and prevent scaling and corrosion on the nozzles. The spray is commanded to run when the IT inlet temperature sensor is healthy and no spray alarm is presented.

4. AHU HE bypass

When the PV system is chosen to controlling the IT inlet temperature, a bypass valve will modulate open when cooling demand is low. This is to reduce the cooling effect of the AHU HE. Again, the bypass valve controls to the same PV as the AHU fan and spray. As with the IT bypass valve, an alarm will be generated to warn when the AHU HE bypass valve is open. As mentioned earlier, the bypass valve is commanded to open when the IT inlet temperature sensor is healthy and the PV system is selected to control the IT inlet temperature.

3.8 Legionella considerations

The bacteria in general and specifically the bacterium Legionella can growth in some equipment that hosts wet environments such as cooling towers and condensers [140]. The Legionella's disease can be contracted by inhaling aerosols containing the bacteria. The risk of growing the bacteria increases if the host temperature is between $25℃$ and $45℃$ and the water is stored or

recirculated for a long time. Thus, the environment of the designed AHU can grow the bacteria and cause potential risks for users and public.

To eliminate the risk of growing the bacteria, some points and procedures are addressed and followed in the design of the AHU:

1. Design the spray system perfectly to ensure that the release of water droplet from the nozzles is properly controlled. The design involves a technique to flush the spray tray automatically every time the spray system is not in use.

2. Ensure that the water cannot stagnate anywhere in the AHU. This is achieved by properly designing the sumps of the AHU and make drain holes in each part of the unit.

3. Avoid using materials that harbour bacteria and other microorganisms in the AHU design such as the wetted media and using direct spray instead.

3.9 Commissioning and start-up of the experimental rig

Commissioning and start-up is the final step before running the experiments. With the AHU, rear door, and the CHx40 installed, it is necessary to start up, test, and commission the facility to ensure that the system performs per the intent and design of the present study to achieve the cooling requirement in this study and for more general data centre cooling applications.

After installing of the AHU and connecting it with the loop of the rear door and the CHx40, the unit is connected to the mains in the control panel of the building, the equipment was started up. The start-up process involved filling the loop with the coolant (which is a mixture of potable water and propylene glycol), check all the valves are opened including the supply water valve of the spray, check the piping loop is leak free especially in the joint points, and verifying that the fan is rotating in the proper direction.

The test procedure of the AHU involved running the fan with variable speed and adjust the damper in the exhaust duct to adjust the air flow. This also

included running the circulation pump and make sure the air is fully purged from the loop to avoid cavitation on the pump and to obtain stable flow rate in the loop.

Commissioning the designed system is the last step before running the system. It includes verifying that the designed cooling system satisfies the cooling requirement of the IT, and the system installed properly to perform per the intent of the design. The commissioning process also includes calibrating the sensors and measurements points as it will be explained in detail in the next chapter.

3.10 Flow issues and designing reconsiderations

In this section, the problems encountered during the operation of the system as well as the solution procedures are identified. The investigations showed some shortcomings in the design and identified some weaknesses that cause the problems in the cooling loops therefore future designs and optimisation procedures should focus on these elements.

3.10.1   Secondary loop problems

The flow in the secondary loop of the system encountered a series of problems which caused the flow to deteriorate and in turn decrease the cooling efficiency of the loop. The secondary loop flow rate history over a period of 16 months of operation is shown in figure 3.25. After commissioning and starting up the system, the flow rate was recorded at about 14 l/min (at maximum pumps speed) which is the designed flow rate of the CHx40 in the current configuration of the servers' modules. However, three months later, the flow rates started dropping as shown in figure 3.25. The flow deterioration continued for over a year until it reached a flow rate of 4 l/min where some of the servers encountered overheating and in some cases reaching the critical temperature for the CPUs and servers shutdown in many cases.
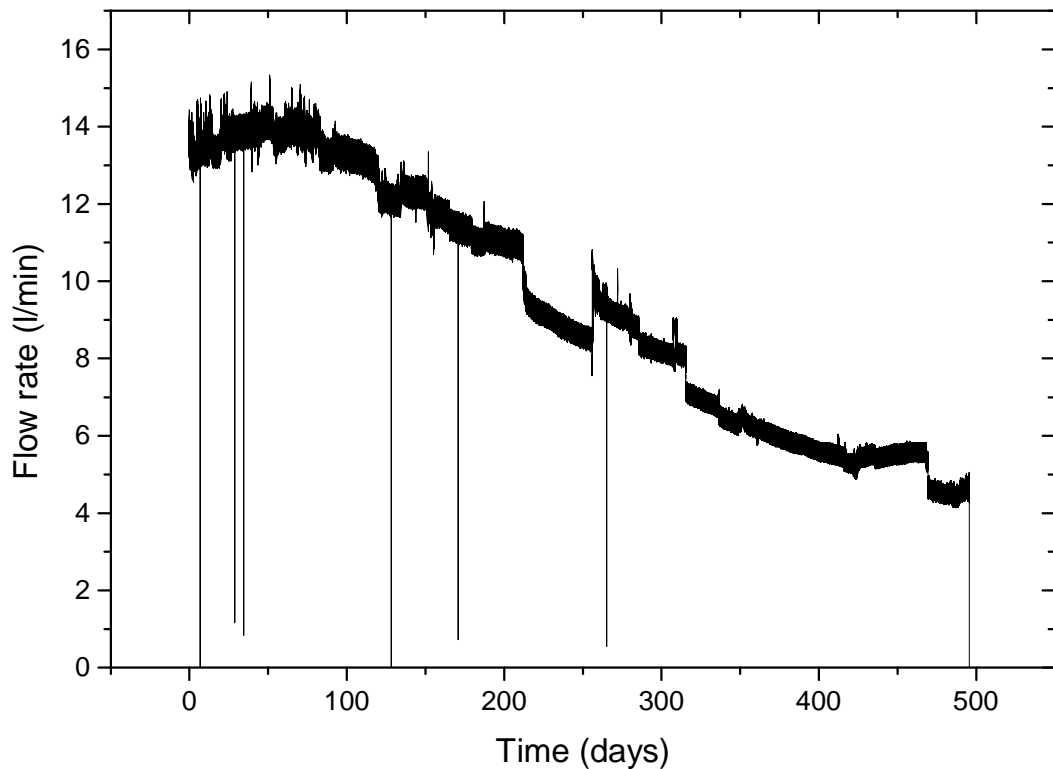
Figure 3.25 History of the deterioration of the secondary loop flow rate.

The issue became very critical and the flow rate diminished further when the module servers' pumps were powered for the first time. At that stage, some of the servers encountered shutdown even when the rack was in an idle load. Thus, a deep investigation of the problem was necessarily to fix the issue and put the system back into operation.

3.10.2 Servers' modules (DCLC heads) issues

The investigations have passed through a series of steps included testing every element in the loop which will be explained in the following points:

1. The first investigation was to check the servers' modules by replacing a selection of them with new ones. This step has shown some improvement in the CPUs temperature of the replaced module's servers which gives an indication that there is deterioration in the servers' modules. However, this was not the only problem as there is no flow in the loop.

2. Before dismounting all the servers' modules from the servers, which is a very delicate process as it can damage the CPUs' pins, there was also a debate that there is air in the loop which could cause the problem. So the loop was pumped with water and the air was released from the pressure relief valve in the top of the loop. However, no improvement was observed and the flow was still very low (less than 1 l/min).

3. The third step included taking one of the servers that encountered overheating out of the rack and dismounting the module server and draining the coolant for investigation. The coolant showed a murky colour and filtration of the sample showed a mixture of particles including oxide products as shown in figure 3.26.
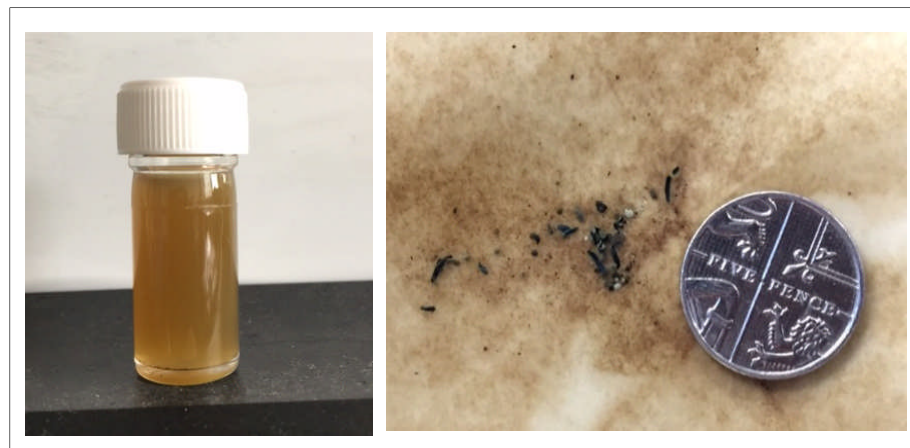
Figure 3.26 Photo of the coolant sample and debris appearing after filtration.

4. The acidity of the coolant was measured and showed a pH of 5. The pH of pure glycol or glycol water solution is approximately in the range between 7 and 8 [141, 142]. The drop in pH has been associated with degradation of the glycol to a mix of, oxalic, lactic, formic and acetic acids. The decrease of pH in the coolant, which is glycol/water mixture, causes the formation of organic acids during thermal oxidation which accelerates the corrosion of metallic components in the loop. As the cold plates are made of copper and the presence of organic acids is

resulted in etching the copper which is likely responsible for the build-up of sludge and potentially failure of other components.

5. Disassembling the server module was necessary at this stage for further investigation of the problem. So the components of the server module were removed and showed build-up of sludge within the unit as shown in figure 3.27. After cleaning the cold plate with water, the plate appeared etched indicating corrosion of copper and the corrosion was also observed in the pipe fixtures.

(a)

(b)

(c)

(d)
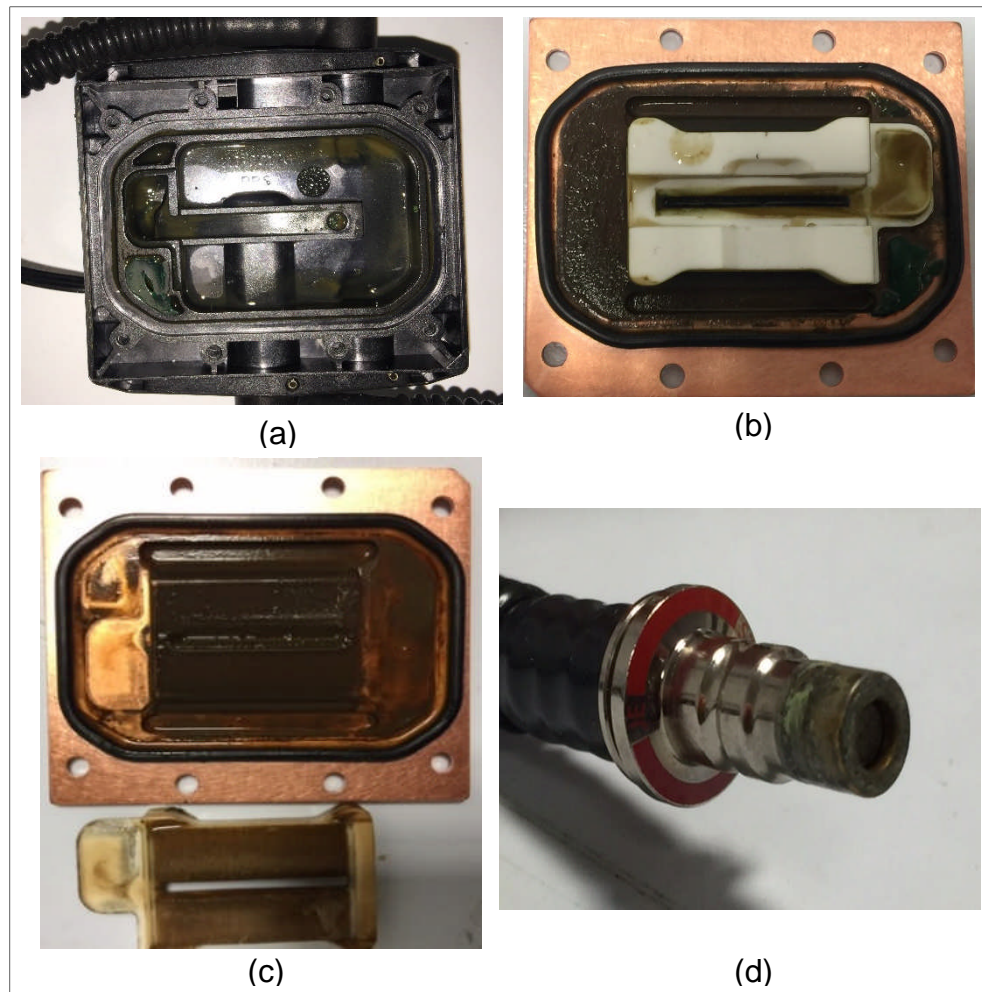
Figure 3.27 photos showing the degradation in the server module head; (a) Shows the casing of the pump, (b) shows the cold plate with the rubber distributer guide, (c) shows the cold plate with the rubber guide removed from its position, and (d) shows stäubli valve degradation.

6. Based on the results of the investigations on the single server module, all the servers had to be taken out of the rack to dismount the servers' modules. Then, each server module was connected separately to a flow loop to check the flow resistance and pressure drop of the module. The results showed that 20 of 30 servers' modules have very high resistance to the flow. These modules were disassembled to be cleaned. The electronic circuits of the pumps were dismantled to keep them away from the water in the cleaning processes.

7. All the cold plates were cleaned in a multistep process. The first step including washing them with water to remove the debris and particles. The second step was to leave them in a weak phosphoric acid in a sonication bath of 90℃ in five stages and each stage includes one hour in the sonication pool. The last step was to wash them with deionised water and leave them to dry. Figure 3.28 shows a photo of a cold plate after cleaning. The pumps casing, gaskets, and the rubber distributer guides were also cleaned with water and left to dry.
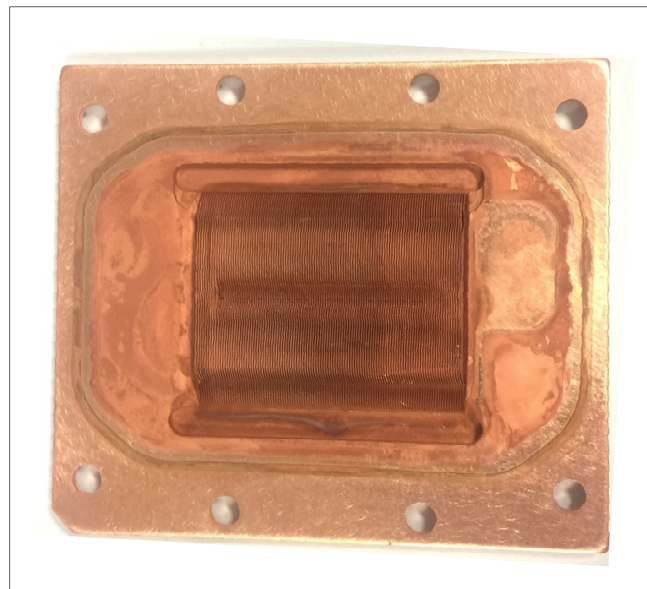
Figure 3.28 Cold plate after cleaning in the sonication path.

8. The final stage before mounting the servers' modules was to assemble all of the components and fill all the servers' modules with the coolant. The process included filling every server module separately using a

small positive displacement pump. The coolant is a mixture of 25% ethylene glycol and 75% deionised water.

9. The manifold and the CHx40 were also dismounted from the rack, drained and cleaned. After that, they were primed by filling them with the coolant mixture to purge the air.

10. The final step at this point was to connect the servers' modules in the manifolds and connect the CHx40 with the loop to start up the loop before mounting the servers' modules in the servers. However, running the system showed a flow rate of 2 l/min which is still far below the required flow rate to keep the servers in a safe thermal environment. Thus further investigation was required to improve the flow rate and solve the remaining problems in the loop.

3.10.3   The CHx40 issues

The CHx40 is the next stage that need investigation to solve the flow problem in the secondary loop. Analysing every components in the CHx40 is necessary to eliminate the number of variables that cause the problem. Before starting the disassembly of the CHx40 components, it is useful to have a general overview on the system by looking at the more obvious items that have the potential to cause the problem such as filters or flow meters (which if broken can cause restriction to the flow). A general inspection showed that the flow meter is spinning freely and does not cause any restrictions. Also no obvious filters were found within the system. Thus, at this stage, an inspection of every item in the CHx40 is required to find the source of the problem. For this, the following steps were followed:

1. To eliminate the interference with the other components and to make sure that the problem is in the CHx40 and not the manifold, the CHx40 was connected in a direct loop to itself as shown in figure 3.29. After priming the loop, the CHx40 was switched on and the flowmeter recorded a flow rate of 2l/min initially and then dropped to 0.5l/min. This

results proved that the problem is in the CHx40 as it recorded very high pressure drop.

2. To eliminate the possibility of the heat exchanger blockage, the heat exchanger was taken off from the loop and the CHx40 showed the same flow rate which means that the problem is either in the pumps or the reservoir.

Figure 3.29 A photo showing the CHx40 where the reservoir is placed in a higher level than the pumps and CHx40 is connected via a hose to itself.

3. The next step was to replace the pumps. The two pumps were replaced with new pumps of the same type as the original ones (D5vario-38/700BM) but with the ability of changing the pump speed to allow studying the effect of the secondary loop flow rates on the cooling performance. However, even when the pumps were replaced, the flow rate stayed low of initially 2l/min and then dropped to about 0.5 l/min as previously.

4. The last component that needed investigation was the reservoir. The reservoir is a sealed container that stores the coolant before being sucked by the pumps as it also works as a pressure vessel to regulate the increase of the pressure resulting from the thermal expansion of

the coolant as well as it has a pressure relief valve that purges air from the loop. However, it was out of the suspicion as it is not expected that a sealed reservoir can contain any restrictions for the flow. By using an industrial laproscope, a blocked filter was found inside the sealed (welded) reservoir as shown in figure 3.30. For this, a hole of 45mm is drilled in the top surface to remove the filter. A loop was formed between the reservoir and pumps to check the flow rate as shown in figure 3.31 and the flow rate was reported to be 23 l/min for the maximum speed of the pumps which are connected in series.
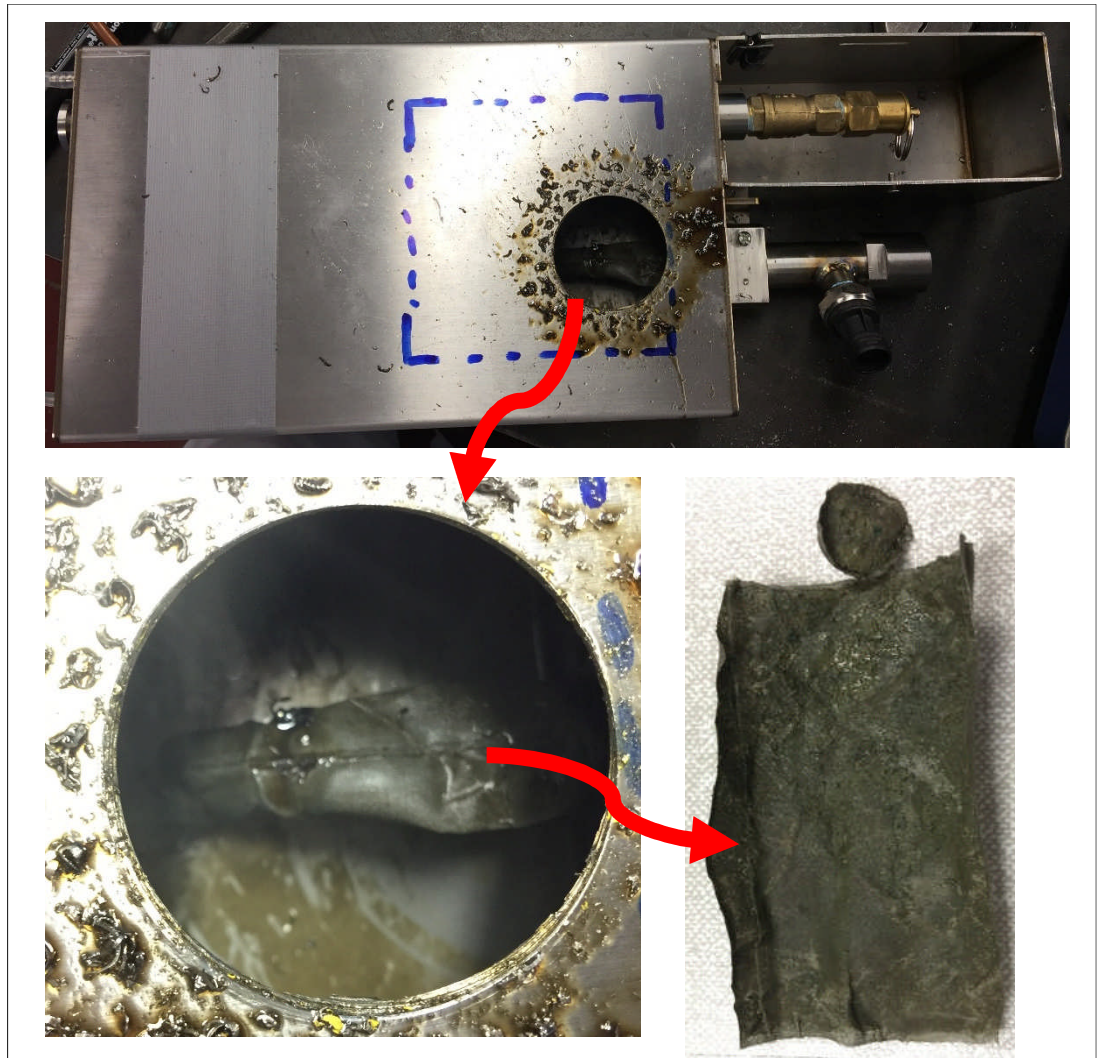
Figure 3.30 The blocked hidden filter in the sealed reservoir after making the hole.

5. The reservoir and the HE were washed with acetic acid solutions and then with water before mounted on the CHx40 chassis. It is worth mentioning that the reservoir was lifted from its original position to be higher than the pumps so that any accumulative air will be in the top of the reservoir as shown in figure 3.32. Then the CHx40 and the manifold were refilled with the coolant mixture. The server modules were mounted on the CPUs of the servers which are separated by a thin layer of thermal paste for better thermal conduction. The servers, manifolds, and the CHx40 were mounted back in the rack, and commissioned to be in operation again with a flow rate of 14 l/min when the CHx40 pumps are only in operation and 16l/min when assisted with the servers' modules pumps. The manifolds were positioned in a lower level than the CHx40 in the new configuration to force the air to be accumulated in the highest point in the loop which is the reservoirs now.
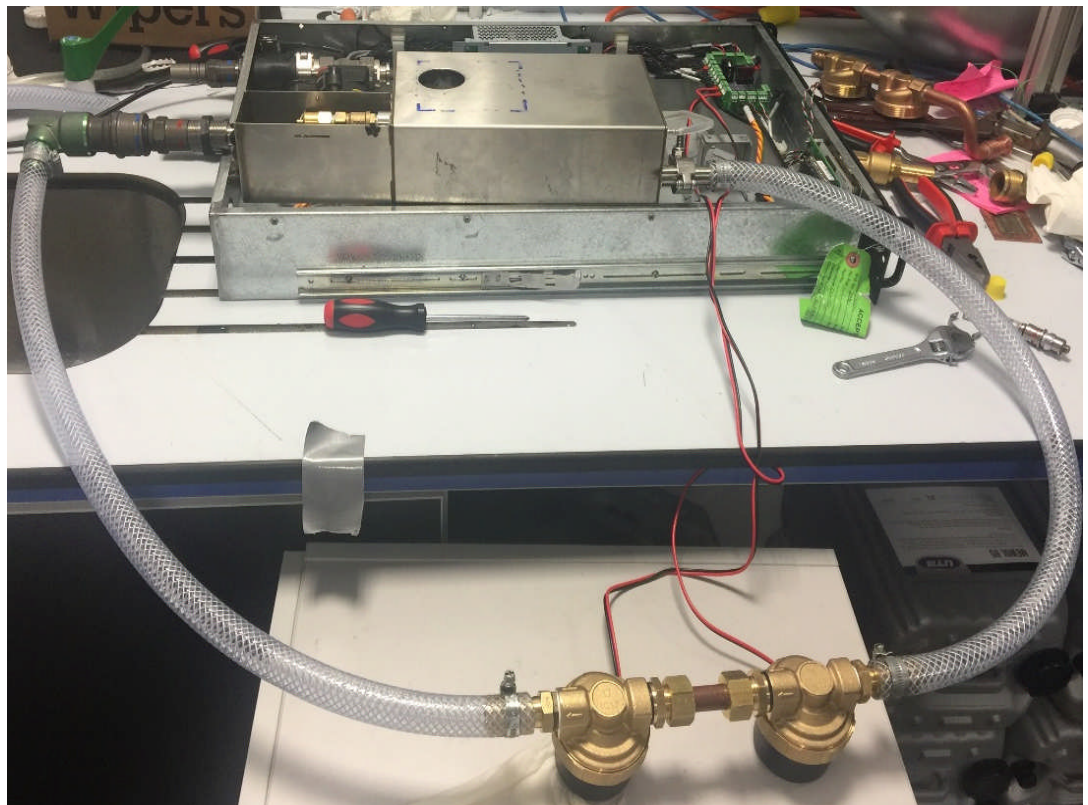
Figure 3.31 The reservoir and pumps loop after removing the hidden filter.
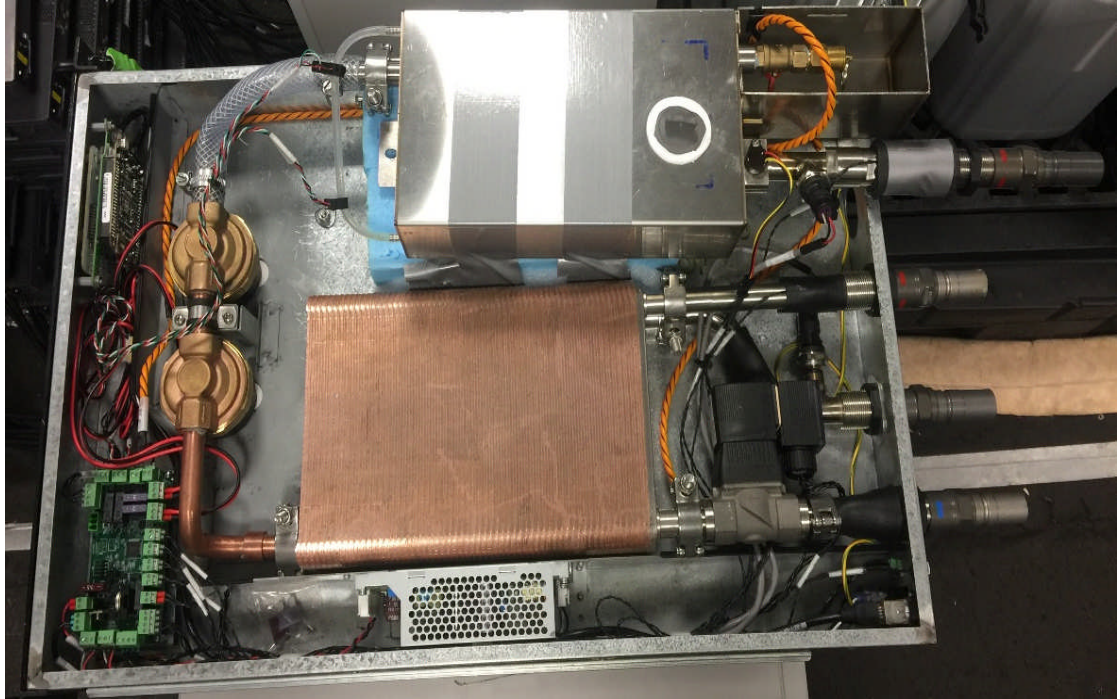
Figure 3.32 The CHx40 design after modification.

3.11 Summary

In this chapter, the design methodology of the direct liquid cooled data centre from the chip level to the facility level has been discussed. This included a full description of each item in the system and the selection of the different components for efficient chiller-less liquid cooled data centres. The control methodology was also presented and discussed as well as the commissioning of the system. Design considerations and shortcomings were highlighted particularly in the secondary loop of the servers where a drop in the flow rate was observed over a lengthy operation time resulted in servers overheating. Therefore, a full systematic investigation methodology to find the source (or sources) of the problem was developed and proposed. Based on this, design modifications were suggested and applied in the current design.

Chapter 4 Experimental setup and test procedures

4.1 Introduction

To reiterate, a dual loop DCLC chiller-less data centre test facility, as shown schematically in figure 4.1, was designed and constructed in the present study to cool a rack of thirty Sun Fire V20z servers consuming approximately 6.5kW of power. The server rack, CHx40, and the rear door HE are located inside the lab and the AHU is located in the plant room which is open to the environment.

The cooling facility consists of two main loops: the primary loop and the secondary loop, as shown in figure 4.1. The purpose of the primary loop is to carry the heat from the IT and to reject it to the external environment outside the data centre room. The heat taken by the primary loop, is mainly transferred by one of two means: the rear door HE and the CHx40. The purpose of the secondary loop is to carry the heat generated by the CPUs inside the servers via the DCLC plates and transfer it to the primary loop across the CHx40. As mentioned earlier, the heat generated by the remaining components of the servers are transferred by the circulated air, driven by server fans only, within the rack which is cooled by the primary water loop using the air-to-liquid rear door HE mounted on the back of the rack enclosure. The designed cooling facility reduces the need for specialised CRAH units for room cooling since a large fraction of the heat generated at the IT is absorbed directly by the secondary loop via the cold plates and indirectly via the rear door HE by the recirculated air.

The coolant in the primary loop does not need to meet the same coolant quality standards desired for use in the secondary loop. The coolant in the secondary loop is ethylene glycol (25%) and deionised water (75%) mixture, while the primary loop is filled with propylene glycol (25%) and potable water (75%). The closed primary and secondary loops require no additional make-up water. The current design also provides easier control of the local rack environment as it uses the two coolant loops and physical isolation of the IT from the external environment.

4.2 Sensors and measurement points

Temperature, flow, pressure, and humidity sensors are located at various locations including the inlet and outlet of every component from the servers to the AHU, including the DCLC cold plates, server CPUs, RAM, rear door HE, CHx40 and the AHU, as shown in figure 4.1. In the present test facility, the location of some of the temperature sensors measure the outlet of a component which also represent the inlet of another component.

For the server level, each server has five temperature measurement points. This includes the two CPUs, RAM, and ambient temperature inside the server. The servers also have tachometers for measuring server fan speed.

The CHx40 connects the secondary loop of the servers and the primary loop with the rear door HE and the AHU. In the secondary loop there are two temperature sensors in the inlet (WT1) and the outlet (WT2), pressure sensors in the inlet (WP1) and the outlet (WP2), and flow meter (M1). The primary loop section of the CHx40 contains two temperature measurement points in the inlet (WT3), and outlet (WT4). The CHx40 also contains relative humidity sensors (RH1), and ambient temperature sensor (AT1).

The rear door HE is connected in the water loop between the AHU and the CHx40 (primary loop of the CHx40). Four temperature measurement points are located for the rear door HE. Two in the water loop at the outlet from the AHU which is considered as the inlet to the rear door HE (WT5) and the inlet temperature to the CHx40 which is considered as the outlet temperature of the rear door HE (WT3). In the air side loop of the rear door HE there are two temperature measurements points one in the inlet (AT2) (which is the exhaust section of the servers) and the other point after the rear door HE (AT3). Each measurement point is an average of multi thermocouple sensors located equally at various sections from the bottom to the top of the rear door.

In the AHU there are three loops as described above, the sensors in the air loop consist of ambient temperature (AT4) and humidity sensors (RH2), post spray temperature sensors (AT5), post HE sensors of temperature (AT6) and relative humidity (RH3), pressure drop sensors across the HE (Δp3) and the fan section (Δp4).

The sensors in the process water loop consist of inlet (WT5) and outlet (WT6) temperature sensors to the IT, inlet (WT7) and outlet (WT8) temperature sensors to the HE, flow meter (M2), and pressure sensors in the inlet (WP5) and outlet (WP6) of the HE.

In the spray water loop, two temperature sensors are installed; a temperature sensor in the spray water (WT9) and temperature sensor in the drain section (WT10). Two flow meters are installed, one in the supply side (M3) and the other one is located in the drain side (M4).

Power measurement sensors are used at various points to measure the energy consumption of the designed facility. At the rack level, power measurements are allocated to measure the power consumption of the servers (PW_Servers). The CHx40 power consumption is measured using a built-in sensor which measures cross power consumption of the CHx40 unit (PW_CHx40). The AHU is powered from the control panel and a power measurement sensor is used to measure the power consumption of the whole unit (PW_AHU).
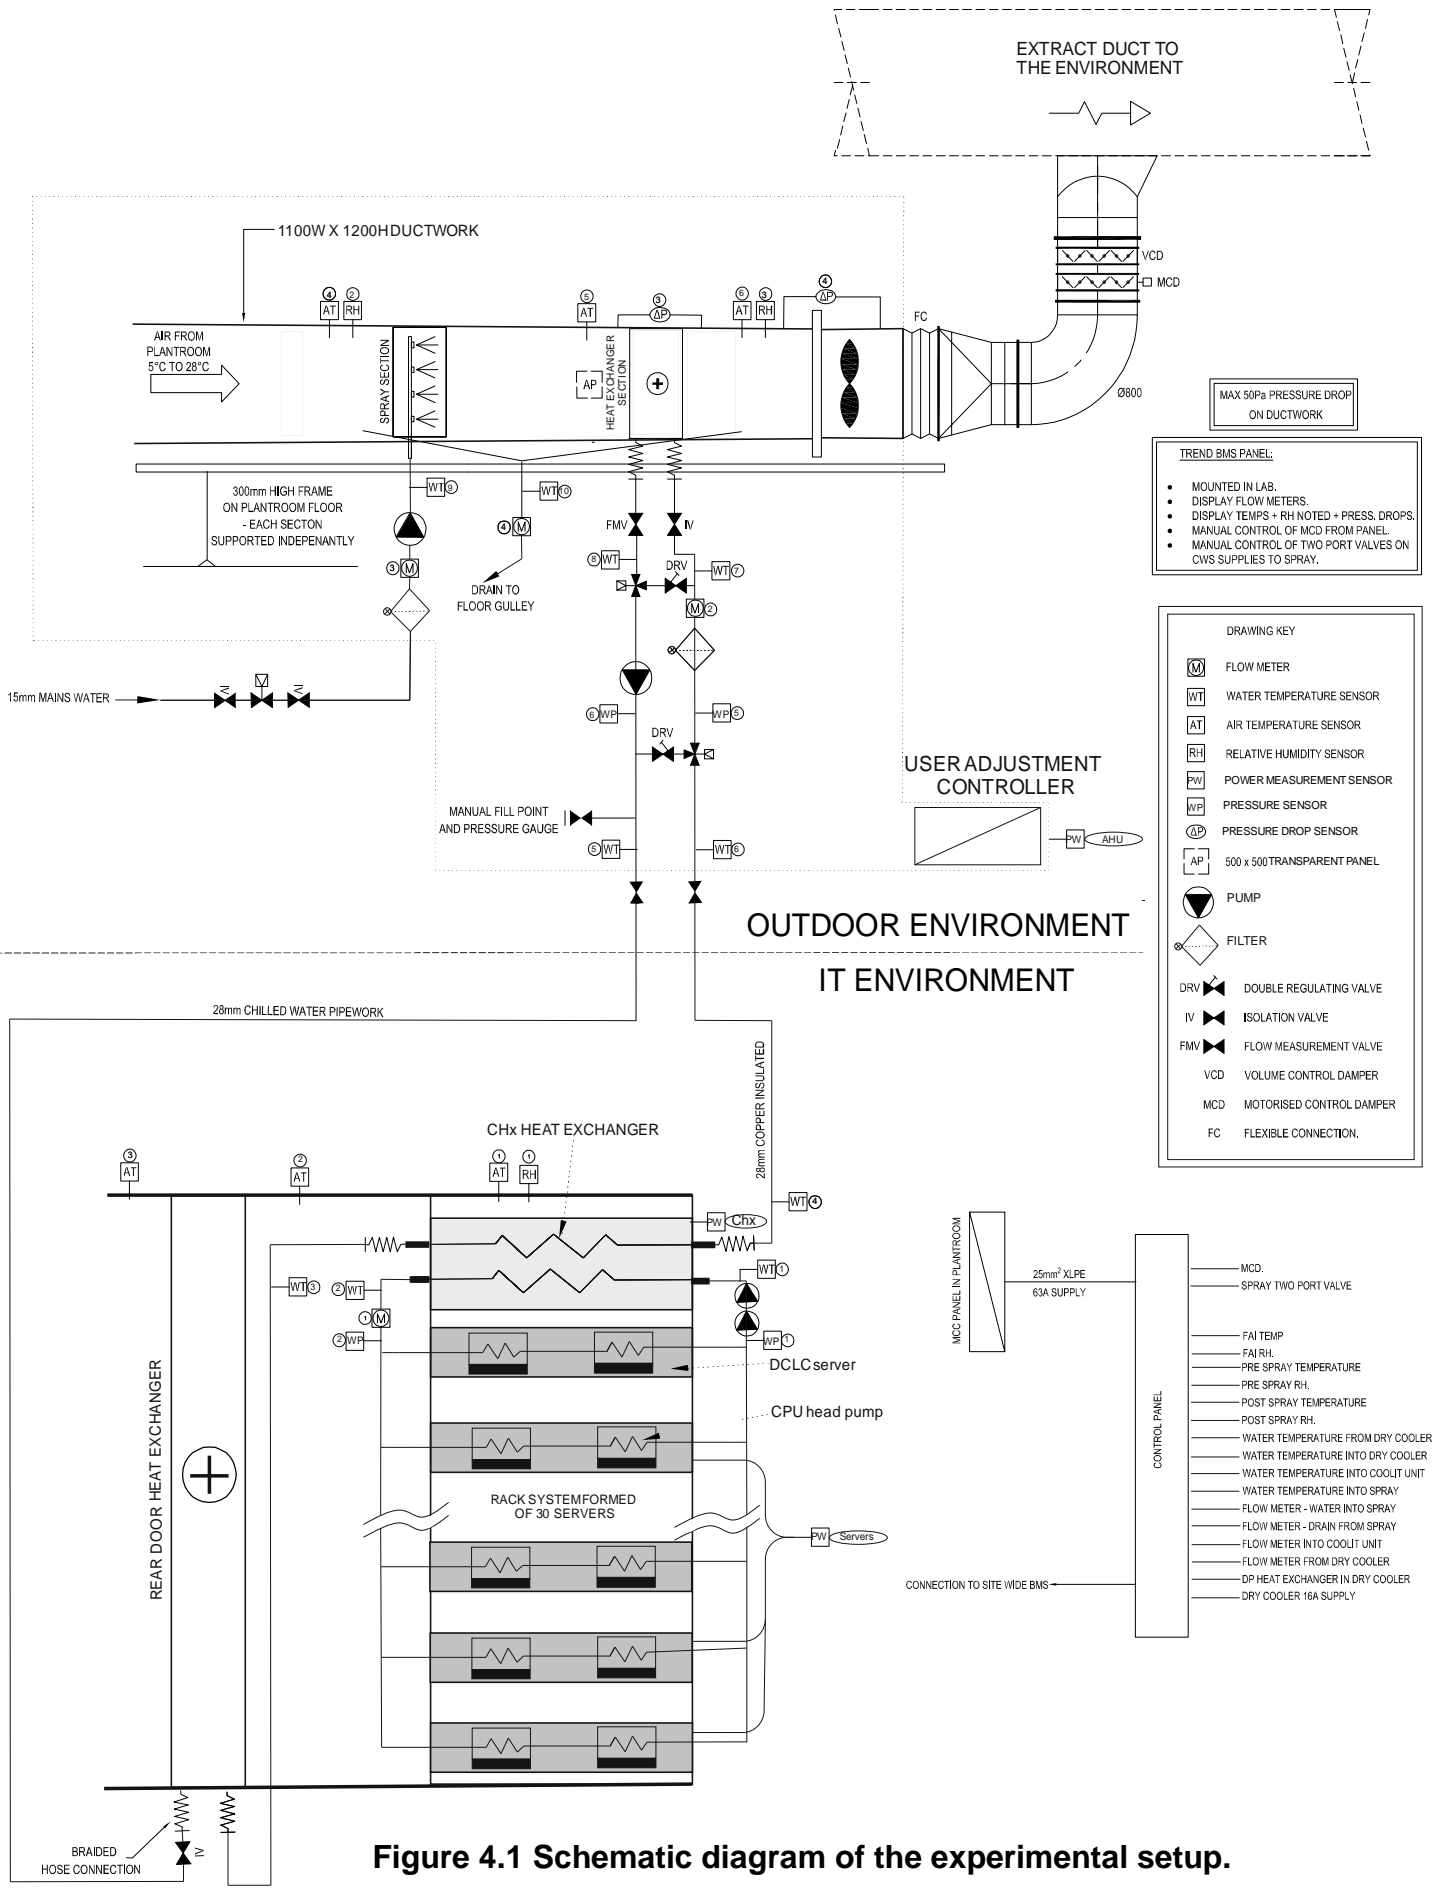
Figure 4.1 Schematic diagram of the experimental setup.

4.3 Sensors calibration and uncertainty analysis

4.3.1 Temperature and humidity measurement

K-type thermocouples were used for all the temperature measurement points of the liquid side of the loops. These thermocouples where all calibrated using a temperature controlled water bath. The calibrations were performed at 5℃ intervals in the temperatures range between 5℃ and 70℃ with the thermocouples resting in the bath for 15 min before 2 min of data was acquired and averaged. A reference thermocouple, which was already calibrated, was also inserted in the bath to give reference temperature readings which have an uncertainty of ± 0.5℃. Once the data was collected, the reference thermocouple and all other thermocouples were calibrated to match its output. The calibration was accomplished through linear functions which describe the offset between the reference thermocouple measurements and all other thermocouple measurements. Since the K-type thermocouples were all calibrated to match the reference thermocouple readings, then the uncertainty in all the temperature readings is equal to the reference thermocouple uncertainty of ± 0.5℃.

Two temperature/humidity (air type) sensors were used in the experiment; one at the inlet section of the AHU (wall type from Carel DPWC110000) to measure the ambient conditions and the other one to measure the air conditions of humidity and temperature (duct type from Carel DPDC210000) after the heat exchanger. The two airside temperature and humidity sensors were calibrated against a reference commercial thermo-hygrometer which is in turn calibrated by the factory (C.A 1244) with an uncertainty of ± 0.3℃ and 1.8% (in the range of 15 to 25℃) for the temperature and relative humidity, respectively. This results in an uncertainty of ± 0.3℃ and 1.8% for the airside temperature and humidity sensors, respectively. An average temperature sensor is also installed after the heat exchanger to ensure accurate measurement of the temperature especially in the wet case operation of the AHU.

4.3.2 Flow rate measurements

Four flow rate measurement points are located in the test facility;

1. Air flow rate in the AHU.

The air flow rate was calculated from the pressure drop across the fan and the fan speed in rpm from the manufacturer fan performance curve. This measurement technique was calibrated against the ASHRAE standard of measuring the air flow rate in ducts [143].

The air flow rate through a duct is calculated from the following equation:

$$\dot{V} = A_d \overline{U} \qquad\qquad \textbf{4.1}$$

Where,

$\dot{V}$ - is the air volume flow rate in m^3/s;

A_d - is the cross-sectional area of the duct in m^2;

$\overline{U}$ – is the average air velocity measured by the traversing method across the cross section of the ducts in m/s.

The average air velocity is measured using the log-Tchebycheff method [143]. This method is based on dividing the duct into concentric circles contain equal points. To obtain the best average velocity, an equal number of readings is recorded from each circular area. As the duct diameter is relatively large (80cm), four concentric circles of 10 measurement points per diameter are chosen. For this, four holes are drilled in the duct at 45° angles from each other. The location of these measurement points is shown in figure 4.2. the arithmetic average of the measured velocities is taken to calculate the mean air velocity:

$$\overline{U} = \frac{\sum_{i=1}^{n} u_i}{n_p} \qquad\qquad 4.2$$

where,

u_i- is the measured velocity in each point (m/s)

n_p- is the number of measurement points $(10 \times 4 = 40)$
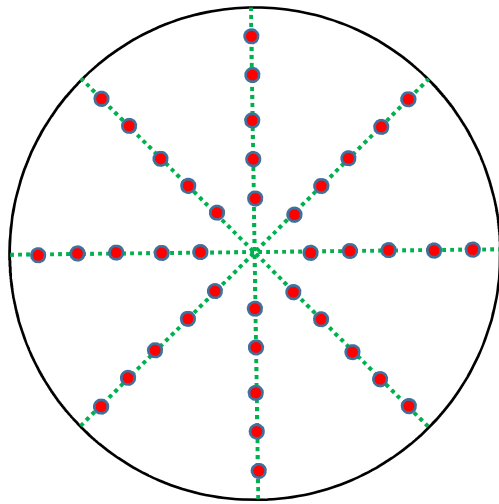
Measurement point distance	Tube Dia	Distance from the wall (cm) for each path
0.019D	80	1.52
0.077D	80	6.16
0.153D	80	12.24
0.217D	80	17.36
0.361D	80	28.88
0.639D	80	51.12
0.783D	80	62.64
0.847D	80	67.76
0.923D	80	73.84
0.981D	80	78.48

Figure 4.2 Locations of measurement points on the cross section of the duct using the traversing method.

2. Process water flowmeter.

The recirculated water flow rate is based on the pressure difference between two known points across the pump. The flowmeter was calibrated in the factory by Airedale, and it was also recalibrated in the lab using a graduated cylinder-stop watch standard. It is worth mentioning that the flow rate increases with stressing the servers as the working temperature is increasing which in turn decrease the density and viscosity of the water.

3. Spray water flowmeter

The spray flow rate is taken from the Carel system. This was obtained based on the characteristics of the pump of the spray pump speed and pressure drop. To ensure that the spray system flow rate is reading correctly, a calibrated external flow meter was installed on the supply section pipe of the spray pump. The latter is only measuring the water consumption, and hence the spray was running on different flow rates and the water consumption was recorded and compared with the spray reading per unit of time.

4. Secondary loop water flowmeter

The flow meter in the secondary loop is of the rotary type and calibrated by CoolIT Systems. The sensor was also recalibrated in the lab using a graded container with stop watch at the ambient temperature of the lab (23℃). The

flow rate in the secondary loop is sensitive to the temperature as the temperature in this loop is higher than the primary loop. It was noted that the flow rate increases by about 1.56 l/min when the rack is stressed at 100% utilisation compared with idle operation.

4.3.3 Pressure drop measurements

The differential pressure transducers were chosen from Johnson Controls (DP2500). The transducers are subjected to a total uncertainty in differential pressure measurement of $\pm 1.5\% + 1Pa$ including general accuracy, temperature drift, hysteresis, linearity, long term stability, and repetition error. The response time of 0.8s is selected using the selectable jumper. The designed operation of the transducer's temperature and humidity ranges from -10 to 50℃ and 0 to 95%, respectively.

4.4 Data logging

All the sensors in the servers and the CHx40 are dynamically logged by the master server (`eng01`) with a network timestamp that can be downloaded and processed later.

The AHU sensors are also dynamically logged with network timestamps and can be accessed via a central programmed panel. The system is remote (over ethernet) monitored and controlled using a pCOWeb card (essentially a linux based ARM system). A HTML page was set up, and hosted by the pCOWeb card which returns valid data from the controller in the pump skid control panel. The card is also configured to send an email with a log of data every 24 hours to a provided email address. A mail server is installed on the headnode (`eng01`) for emails to be forwarded to the email client with the data.

4.5 SPECpower_ssj2008

SPECpower_ssj2008 [144] is a benchmark product proposed by the Standard Performance Evaluation Corporation (SPEC). It is designed to use Java server applications to provide a view of a server system's power consumption. The general approach of the SPEC is to compare measured power

consumption with measured performance. The main requirement of the SPEC is to include power measurement data of a server running at different target load levels which reflects the fact that data centre server system runs at different load targets relative to the greatest throughput. The SPECpower_ssj2008 consists of three software components [145]:

1. Server Side Java (SSJ) – Workload

The SSJ workload is a Java program which has the capability to exercise the CPUs, memory, Java virtual machine (JVM), just in time compilers (JIT), and other aspects of the operating system of the server under test (SUT). The SUT is chosen to be `eng06` in the present rack configuration. The essential function of the SSJ is to control the workload instance, synchronise different workload instances, track the state of the workload, and consolidate the results and pass it to the network (CCS). The SSJ is also responsible for the schedule and executes transactions, determine workload intervals and state, and generate reports.

2. Power and Temperature Daemon (PTDaemon)

The PTDaemon is installed on the control server (`eng01` in the present rack configuration) to offload the work controlling the power analyser during measurement intervals from the SUT (`eng06`) to the controlled servers `eng01`.

3. Control and Collect System (CCS)

The CCS is a multi-threaded Java software which controls and enables the coordinated collection of data between `eng01` and `eng06`.

The basic hardware of the SPECpower_ssj2008 consists of three components that work together to gather the server power consumption and performance data by applying `eng06` with a predefined workload. These components are the SUT, Power Analyser, and Controller system as shown in figure 4.3

111
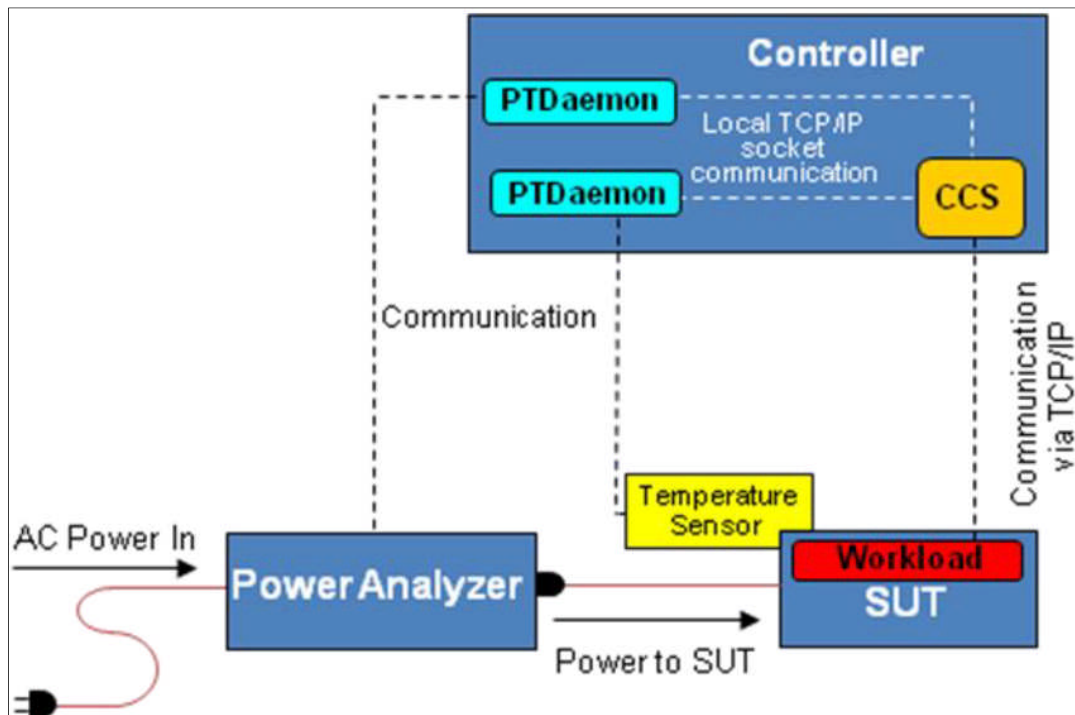
Figure 4.3 Architecture illustration of SPECpower_ssj2008 [145].

The target loads are measured based on the performance of a fixed specified workload that runs on eng06. The completed number of transactions per second over a fixed period which are generated through a mix of transactions by the Java application workload is used to define the server performance. The application starts by conducting a calibration phase used to determine the maximum output of the server, by establishing the maximum rate of transactions which can be completed. The calibration phase consists of three stages, each stage has been chosen to be of five-minute iterations, allowing the Java runtimes the opportunity to utilise just-in-time compilation and other runtime adaptations sufficiently; the maximum output used to determine the target loads is the average of the output realised in the last two iterations in the calibration phase.

After the maximum output has been identified, the Java application then determines the specified output load value which corresponds to a percent of the target load of the maximum value as calibrated (for example 25% load which means 25% of the maximum load). After this is obtained, the benchmark then enters the measurement period during which the workload iterates

through the specified target loads where they are specified according to the experiment. When applying more than one load in the same experiment, every single load generating transactions for completion at the specified target rate, providing a randomized distribution to determine the times at which transactions are issued.

A Voltech PM1000+ [146] is used to measure the power during the experiments which is logged and monitored. The Voltech PM1000+ is a powerful and wide application use precision power analyser. The power measurement is based on the basic definition of the true power:

$$P = \frac{1}{T} \int_0^T F\, V_i I_i \, dt \qquad\qquad 4.3$$

Where, P is the power, T is the total time of the experiment, F is the power factor V_i and I_i are the instantaneous voltage and current respectively.

4.6 Synthetic load generation and load scenarios

The rack servers are stressed with a range of loading levels using Stress Linux. Stress Linux is compatible with the command line which makes the process of writing the scripts easy and flexible to achieve the required load target. As mentioned in the previous chapter, the server in the top of the rack (`eng01`) is the master server which is connected to the remaining servers via a local network, the Stress Linux script is implemented through a secure channel. For this, a secure shell (SSH) is provided by the secure channel to send the stress command line to the slave servers (`eng02` to `eng30`). This also enables the master servers to connect to the IPMI of the slave servers for monitoring purposes such as the temperature readings of the embedded thermocouples in the servers.

4.7 Characterisation of the data centre components

Thermal and hydraulic characterisation of liquid cooled data centres is crucial to establish metrics and behaviour of the system for various alterations of the operating conditions within the cooling loops. The thermal resistance across

heat exchange stages in the data centre cooling systems is a key interest point to data centre designers. The comprehensive thermal resistance of the data centre from the heat generation to the environment will be determined by characterising these individual thermal resistances. Under varying conditions, it is important to ensure thermal benefits are derived from particular operating conditions which can be achieved by measuring power consumption of the cooling equipment to make sure it does not come with a significant energy cost.

In the present work, the three main heat exchange units in the cooling system will be characterised regarding the effect of varying the secondary and primary flow rates of each unit. To represent the steady state condition of the system under operational conditions, each test is carried out for 40 minutes, and the average of the last five minutes of the collected data of each test is used as steady state conditions are expected to have been reached.

4.7.1 CHx40 unit

The impact of the flow variation in the internal and external loops on the thermal resistance of the CHx40 heat exchanger is investigated. The approach thermal resistance of the CHx40 (R_{CHx}), equation 4.4 is the relevant thermal metrics of significance:

$$R_{CHx} = \frac{WT_1 - WT_3}{Q_{CHx}} \qquad\qquad 4.4$$

Where, Q_{CHx} is the heat exchange across the CHx40, calculated from the average between the sensible heating of the primary water loop and the sensible cooling of the secondary water loop as they pass through the CHx40 unit.

$$Q_{CHx} = \frac{Q_{P_{CHx}} + Q_{S_{CHx}}}{2} \qquad\qquad 4.5$$

$$Q_{P_{CHx}} = \dot{m}_2 C_{pw}(WT_4 - WT_3) \qquad\qquad 4.6$$

$$Q_{S_{CHx}} = \dot{m}_1 C_{pw}(WT_1 - WT_2) \qquad\qquad 4.7$$

The experiments are conducted at constant inlet temperature (25℃) to the CHx40 in the primary loop. This will be controlled by setting the AHU in auto

spray mode and selecting the inlet temperature to the IT as the set point variable. The rack is loaded at full power (100% utilisation) for all the servers. The secondary loop flow rate is altered manually using the two speed controllers of the two pumps in the CHx40. The primary loop flow rate is changed using the controller from the pCOweb page.

4.7.2 Rear door unit

The RDHx thermal resistance is altered by the small fans of the servers which are PID controlled depending on the stress load and the primary flow rate where the water is passing through the rear door before going to the CHx40 HE. For this the RDHx thermal resistance will be calculated for different primary loop flow rate where it has a direct effect on the performance while the effect of the AHU fan speed will also alter the heat rejected from the loop and will also be investigated.

The RDHx thermal resistance can be calculated from equation 4.8. The actual heat exchange to the RDHx will be calculated based on the water side as the associated error will be high if the heat transfer is calculated from the air side of the heat exchanger.

$$R_{RDHx} = \frac{AT_2 - WT_5}{Q_{RDHx}} \qquad\qquad 4.8$$

$$Q_{RDHx} = \dot{m}_2 C_{pw}(WT_3 - WT_5) \qquad\qquad 4.9$$

4.7.3 AHU unit

The thermal resistance of the AHU can be characterised based on the primary water flow rates and the air flow rates in the heat exchanger. The thermal resistance approach for the AHU heat exchanger (R_{AHU}) is given in equation 4.10.

$$R_{AHU} = \frac{WT_7 - AT_4}{Q_{AHU}} \qquad\qquad 4.10$$

The heat exchange through the AHU heat exchanger (Q_{AHU}) is calculated from the average heat gained by the passing air (Q_a) (the air and spray in the spray case) and the primary water ($Q_{P_{AHU}}$) that passes through the heat exchanger.

$$Q_{AHU} = \frac{Q_{P_{AHU}} + Q_a}{2}$$ 4.11

Depending on the setting of the AHU operation, the spray system can be on or off. It is also possible to have both scenarios in the same experiment depending on the weather and setting conditions. For both dry and spray cases, the heat transfer on the water side is calculated using equation 4.12.

$$Q_{P_{AHU}} = \dot{m}_2 C_{pw}(WT_7 - WT_8)$$ 4.12

In the dry case, the air side capacity is calculated using equation 4.13

$$Q_a = \dot{m}_a C_{pa}(AT_6 - AT_4)$$ 4.13

For the spray case, the heat transferred in the air side is the sum of the heat transfer to the passing air ($Q_{a,s}$) which includes sensible and latent heat exchange, and the sensible heat transferred to the spray water (Q_{sw}).

$$Q_a = Q_{a,s} + Q_{sw}$$ 4.14

$$Q_{a,s} = \dot{m}_a(h_{a,out} - h_{a,in})$$ 4.15

The latent heat $h_{a,out}$ is at AT_6 and RH_3, and the $h_{a,in}$ is at AT_4 and RH_2 conditions.

$$Q_{sw} = Cp_w \dot{m}_r(WT_{10} - WT_9)$$ 4.16

$$\dot{m}_r = (\dot{m}_3 - (x_{a,out} - x_{a,in})\dot{m}_a)$$ 4.17

The humidity ratio $x_{a,out}$ is at AT_6 and RH_3, and the $x_{a,in}$ is at AT_4 and RH_2 conditions.

Several performance parameters are also introduced to characterise the spray cooling conditions beside the thermal resistance and the effectiveness. Spray efficiency (SE) is defined as the measure of how much water actually evaporates into the air for different air flow rates and spray flow rates as shown in equation 4.18 which is defined as the ratio of the evaporation rate to the spray flow rate.

$$SE = \frac{\dot{m}_a(x_{a,out} - x_{a,in})}{\dot{m}_3}$$ 4.18

Where $\dot{m}_a$ is the air flow rate in kg/s.

The pressure drop across the heat exchanger is also important as it is associated with the fan power and in turn the total power consumption of the data centre. For this, the pressure drop penalty ratio (Δp_r) is also calculated to identify the effect of the air and spray flow rates on the pressure drop across the heat exchanger. The pressure drop penalty ratio is defined as the ratio between the pressure drop in the spray conditions (Δp_{spray}) at a specific spray flow rate to the pressure drop in the dry conditions (Δp_{dry}).

$$\Delta p_r = \frac{\Delta p_{spray}}{\Delta p_{dry}}$$

4.19

4.8 Impact of weather and data centre operating conditions in long duration operation

To study the impact of weather conditions and mimic the different operation loads on the thermal performance and cooling energy consumption of the whole data centre, the data centre facility is tested in different weather conditions under the following data centre loads:

1. One hour comprised of 100% stress pulses for 5 minutes and another 5 minutes of idle.

2. One hour comprised of 100% stress pulses for 15 minutes and another 15 minutes of idle.

3. One hour of Idle operation

4. One hour of 25% utilisation.

5. One hour of 50% utilisation

6. One hour of 75% utilisation

7. One hour of 100% utilisation

This load scenario is proposed to mimic a potentially real dynamic operational workload of a real data centre. The set point temperature was maintained

constant during the 7 hours of experiments so that the AHU has to regulate the temperature of the data center with both modes of dry and wet conditions.

4.9 Performance metrics of data centre

Many metrics have been established by research organisations and various industrialists to assess the performance of data centres. As mentioned in chapter two, there are more than thirty metrics that have been proposed in the literature to determine different aspect of the efficiency of data centres. However, the power usage effectiveness (PUE) has been chosen as the standard metric to evaluate the effectiveness. The PUE is a simple metric which provides an easy way to understand the data centre infrastructure energy consumption for a specific IT load. The PUE is defined as the ratio of the total power consumed by the whole data centre facility to the power consumed by the IT equipment as in equation 4.20. The power consumed by the facility includes, the IT, lighting, cooling infrastructure, fire suppression systems, UPS, etc [32].

$$PUE = \frac{P_{facility}}{P_{IT}} \qquad\qquad 4.20$$

In this present work, the cooling load and the IT load will be considered as the facility power consumption only. For this, the partial power usage effectiveness (pPUE) is defined instead of the PUE.

$$pPUE = \frac{P_{IT} + P_{cooling}}{P_{IT}} = 1 + \frac{P_{cooling}}{P_{IT}} \qquad\qquad 4.21$$

Where the $P_{cooling}$, is the total cooling power consumed by the CHx40 (P_{CHx}) and the AHU (P_{AHU}).

$$P_{cooling} = P_{AHU} + P_{CHx} \qquad\qquad 4.22$$

$$P_{AHU} = P_{fan} + P_{process\ pump} + P_{spray\ pump} + P_{control\ unit} \qquad\qquad 4.23$$

$$P_{CHx} = 2P_{CHx\ pump} + P_{CHx\ control\ unit} \qquad\qquad 4.24$$

In practice, there is a variability in defining the borders between the IT and the cooling side in calculating the effectiveness of data centres. Although, the PUE is widely used in considering the performance of data centres, but it has

shortcomings regarding not being able to account for the cooling and the power distribution losses inside the IT. For this, the total usage effectiveness (TUE), which was introduced by [38], will be used. The TUE is defined as the ratio of the total energy consumed by the whole facility of the data centre to the power supplied to the computational components inside the IT ($P_{compute}$).

$$TUE = \frac{P_{facility}}{P_{compute}} \qquad\qquad 4.25$$

In the present study, it is difficult to measure the power losses in the power supply unit (PSU) and the voltage regulator (VR) of the servers. Thus, the power losses in the PSU and the VR will be neglected. As a result, a new metric called total cooling usage effectiveness (TCUE) is proposed. The TCUE metric can be defined as the ratio of the total cooling power consumption in the data centre including any cooling devices inside the servers to the power supplied to the computational components inside the IT. The power consumption of the fans inside the servers will be calculated from the fan speed-power consumption curve that was obtained by [115] in a study on a single Sun Fire V20z server. As mentioned earlier, every server in the rack has four small fans for cooling the PSU and RAM. The server fan speeds are logged and the power consumption of every fan in the servers of the rack for a specific IT operation is captured from the fan curve in figure 4.4.
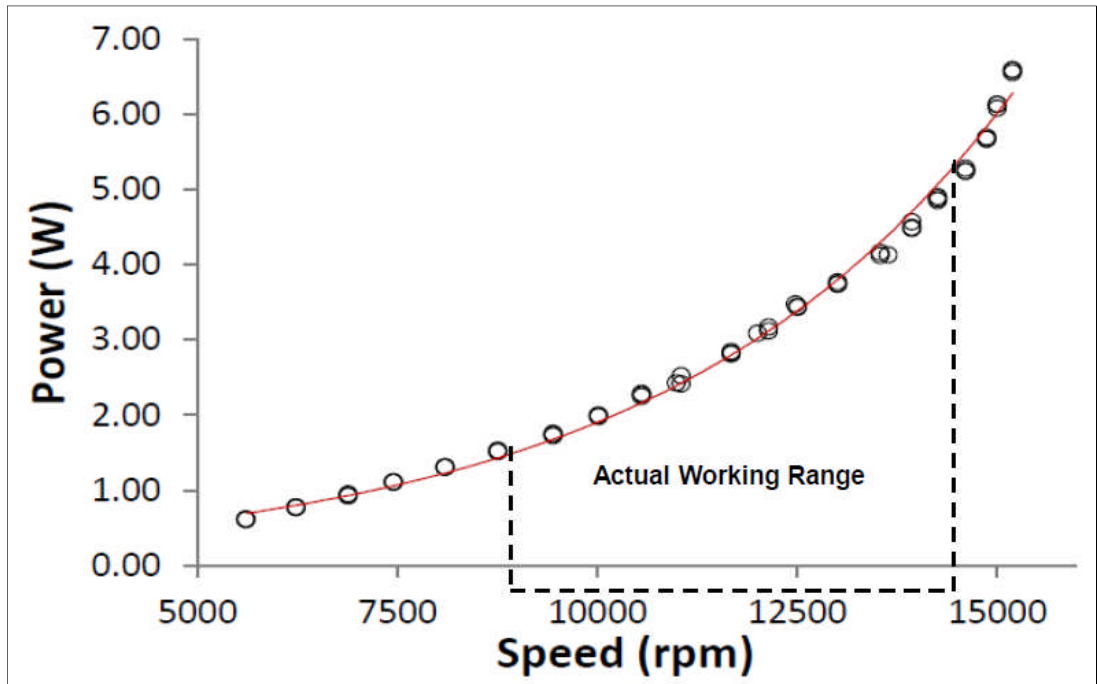
Figure 4.4 Servers small fan speed versus power consumption [115].

As mentioned in chapter three, the CPU pumps can be active or passive depending on the settings of the experiments as will be explained intensively later in the chapter. The CPU pumps can be powered from the server's PCB (printed circuit board), and they all have a fixed speed and power consumption of 1.44W each.

Hence the TCUE is calculated from the following relationship.

$$TCUE = \frac{P_{cooling}}{P_{compute}} \qquad\qquad 4.26$$

$$P_{cooling} = P_{AHU} + P_{CHx} + P_{sc} \qquad\qquad 4.27$$

$$P_{sc} = P_{cpu\ pumps} + P_{servers\ fans} \qquad\qquad 4.28$$

Where $P_{servers\ fans}$ is the total power consumed by the fans of the servers, and $P_{cpu\ pumps}$ is the total power consumed by the DCLC head pumps of the CPUs in the active case.

The $P_{compute}$ will be calculated from the total rack power consumption (P_{rack}) subtracted from the cooling components inside the servers.

$$P_{compute} = P_{rack} - P_{sc} \qquad\qquad 4.29$$

The pPUE and its modified version in the present study $TCUE$, stands for data centre effectiveness not efficiency. Thus, the effectiveness cannot be considered as a comprehensive metric to define the full operational performance of data centre. As a result, a metric based on the energy efficiency of the compute was developed as found in [42]. This metric is called performance per power consumption (performance/Watt). In the present study, the aforementioned SPECPower2008 is used to measure the number of transactions of the CPUs per power consumption of the server in the units of FLOPS/Watt and will be called workload power efficiency (WPE).

$$WPE = \frac{number\ of\ CPUs\ operation\ in\ FLOPS}{average\ power\ consumption\ of\ the\ server} \qquad\qquad 4.30$$

The gap between the effectiveness (pPUE) and efficiency (WPE) of the data centre makes it possible to improve one of them at the expense of the other one. For example, increasing the inlet water temperature to the rack will improve the pPUE as a result of less required power to run the cooling system this has a negative effect on the IT performance as the CPU efficiency reduces with increasing the CPU temperature. For this, a metric was found in the literature, which was introduced by [44], bridging the gap between the effectiveness and efficiency metrics which is called Data Centre Workload Power Efficiency (DWPE). The DWPE is defined as the ratio between the efficiency and effectiveness of data centre. As the pPUE is calculated in the present study and not the full PUE, hence the partial DWPE (pDWPE) will be defined in the present study.

$$pDWPE = \frac{WPE}{pPUE} \qquad\qquad 4.31$$

4.10 Secondary loop pumping configuration and flow distribution

In the process of optimising the data centre operation by using minimal energy consumption per maximal IT performance, the 30 Sun Fire V20z servers are tested in a rack level and the effect of pumping configuration and coolant flow

rate in terms of power consumption, IT performance, and the server component temperatures are reported. The study provides comparative data for the pumping configuration of central and distributed setups. As mentioned previously, each server has two CPUs which are cooled by two DCLC cold plates integrated with two pumps connected in series while the remaining components of the servers are air cooled. The configuration that consists of the DCLC pumps being powered on is referred to as distributed pumping. While, the configuration with the DCLC pumps being powered off and the flow is only pumped by the two central pumps that are connected in series in the CHx40 is termed central pumping. It is important to mention that the DCLC pumps are of a fixed speed and the CHx40 pumps remained the same in the distributed pumping configuration. The flow rate in the secondary loop is controlled manually by changing the CHx40 pumps' speeds.

4.10.1 Test conditions

The primary goal of the experiments is to measure CPUs temperatures and cooling power consumption associated with the servers at different coolant flow rate varying from 4.5 l/min to 13.5 l/min for central pumping and from 7.2 l/min to 15.6 l/min for distributed pumping.

The first part of the experiment is to set the desired inlet temperature to the CHx40 which is regulated by the AHU. This is chosen to be 20℃ to ensure that the inlet temperature to the rack falls within the ASHRAE W4 liquid cooled server envelope [52]. The primary loop flow rate is also kept constant at 18 l/min for all the experiments.

The second part of the experiment is to load the rack with the required synthetic computational load. Five load levels are tested in the experiments which are idle load (0%), 25%, 50%, 75%, and full (100%) utilisation (U).

The third part of the test is to set the secondary loop flow rate. As mentioned earlier, the flow rate is changed manually using the CHx40 pumps speed controller in both configurations of central and distributed. Five speeds are chosen that give a specific flow rate in each case. It is worth mentioning that

both pumps are set to the same speed in each experiment to keep the redundancy of the pumps design.

4.10.2   Measurement methodology

The experimental process for each case is started by setting the inlet temperature to the CHx40, then loading the rack with a specified synthetic computation load. At this point, the chosen secondary loop flow rate is set. The test for each flow rate runs for 40 minutes before changing the pump speed for the next test. Every experiment is repeated three times to check repeatability. The data is collected and averaged for the last 15 minutes as steady state is likely to have been reached. Finally, the steady state data of the three experiments are averaged for further processing.

4.10.3   Flow distribution analysis inside the rack

The role of uniform flow distribution of the coolant in the servers of the rack is very important to ensure the reliability of the servers. Coolant discharge and collection into and from the servers' branches of the distribution and collection manifolds is accompanied by static pressure losses due to the manifold internal wall friction and the change of the coolant momentum as it passes through the branches. In the servers' branches, the pressure falls due to the friction of the wall while the pressure rises due to the sudden changes in the direction of portion of the coolant. Contrarily, the pressure falls in the manifolds due to the direction changes. Thus, the flow rate in the servers' branches vary even for identical DCLC heads as it is impossible to keep the static pressure inside the manifolds constant.

To understand and examine the flow distribution inside the rack, the open source EPANET 2.0 software [147] is used. EPANET allows simulation of water hydraulic behaviour and quality within pipe networks. The network consists of pipes, pipe junctions, storage tanks, pumps, and valves. It calculates the flow rate of water in each pipe in the network, the pressure at each point in the system, the water level in the tanks or reservoirs, water quality and water age calculations. EPANET is based on analytical experiments for major and minor losses in the pipe networks.

The simulation procedures start by drawing the network representation of the secondary loop system. Two configurations are simulated, one for the central pumping and the other for the distributed pumping. The next step, is to define the properties of the nodes, tank, pipes and pumps. This step includes defining the minor losses coefficients, friction factor of the pipes, length of the pipe, diameter, elevation of every node, and pump characteristic curves. The secondary loop conditions are divided into three sections:

1. DCLC heads.

The number of the DCLC heads is 30 which represents the whole rack. Each DCLC head is represented as a pipe of equivalent head losses in the case of central pumping and represented as a two pipes segment connected by two pumps in series in the distributed pumping case, as shown in figure 4.5 and 4.6, respectively. The DCLC head's pump curve is fed into the software in the case of distributed pumping. The minor losses coefficient for the DCLC heads is taken from a recent study by Alkharabsheh et al [92] where they analysed the pressure drop in a direct liquid cooled rack. The DCLC heads are assumed identical in the present analysis. The DCLC heads have a higher pressure drop in the loop as it has corrugated hoses, microchannel heat sinks, change of flow direction and plumbing. Table 4.1 shows the conditions that define the server modules in the EPANET software.

Table 4.1 Defining DCLC head properties in EPANET

Number of server modules	Diameter (mm)	Pipe length (m)	Losses coefficient	Roughness
30	6	2.45	350	0.001524

2. Discharge and collection manifolds

Each manifold is divided into 42 sections (84 sections in total) to represent the Stäubli valve T-junctions. The manifolds have square cross sections, thus the characteristic hydraulic diameter is calculated to be fed into the software. The pressure losses in the manifolds are calculated based on defining the two losses coefficients of major due to the friction and minor losses due to the change in the flow direction and T-junction losses. The minor losses

124

coefficients are fed into the EPANET based on the minor losses coefficient standard tables [148]. Table 4.2 shows the conditions that define the manifolds in the EPANET analysis.

Table 4.2 Defining properties of the two manifolds in EPANET

Number of segments	Hydraulic diameter (mm)	Segment length (mm)	Losses coefficient	Roughness
84	25	45	4	0.072

3. CHx40 module unit

The secondary loop of the CHx40 contains mainly pipes, fittings, Stäubli valves, a tank, two pumps and a heat exchanger loop. The tank and the two pumps will be physically represented in the modelling while the other components will be represented by the head losses. The pump characteristic curves are fed into the program, and every pump has five curves depending on the speed settings used, as shown in figure 4.7. The tank in the CHx40 is of rectangular shape of dimensions of $(30 \times 15 \times 10 \ cm^3)$, hence the tank is fed as a cylindrical tank to the software with an equivalent diameter of 19.53 cm $(D_{eq} = 1.128 \times \sqrt{W \times L})$. Table 4.3 shows the conditions that define the reservoir characteristics in the EPANET analysis.

Table 4.3 Defining properties of the reservoir in EPANET

Equivalent diameter (cm)	Max level (mm)	Min level (mm)	Initial level (mm)	Elevation (m)
19.53	100	0	80	2.1

The software results are validated first against the experimental data and then they will be used to study the flow distribution for the current design, to analyse and then optimise the design for better flow distribution for all the servers in the rack.
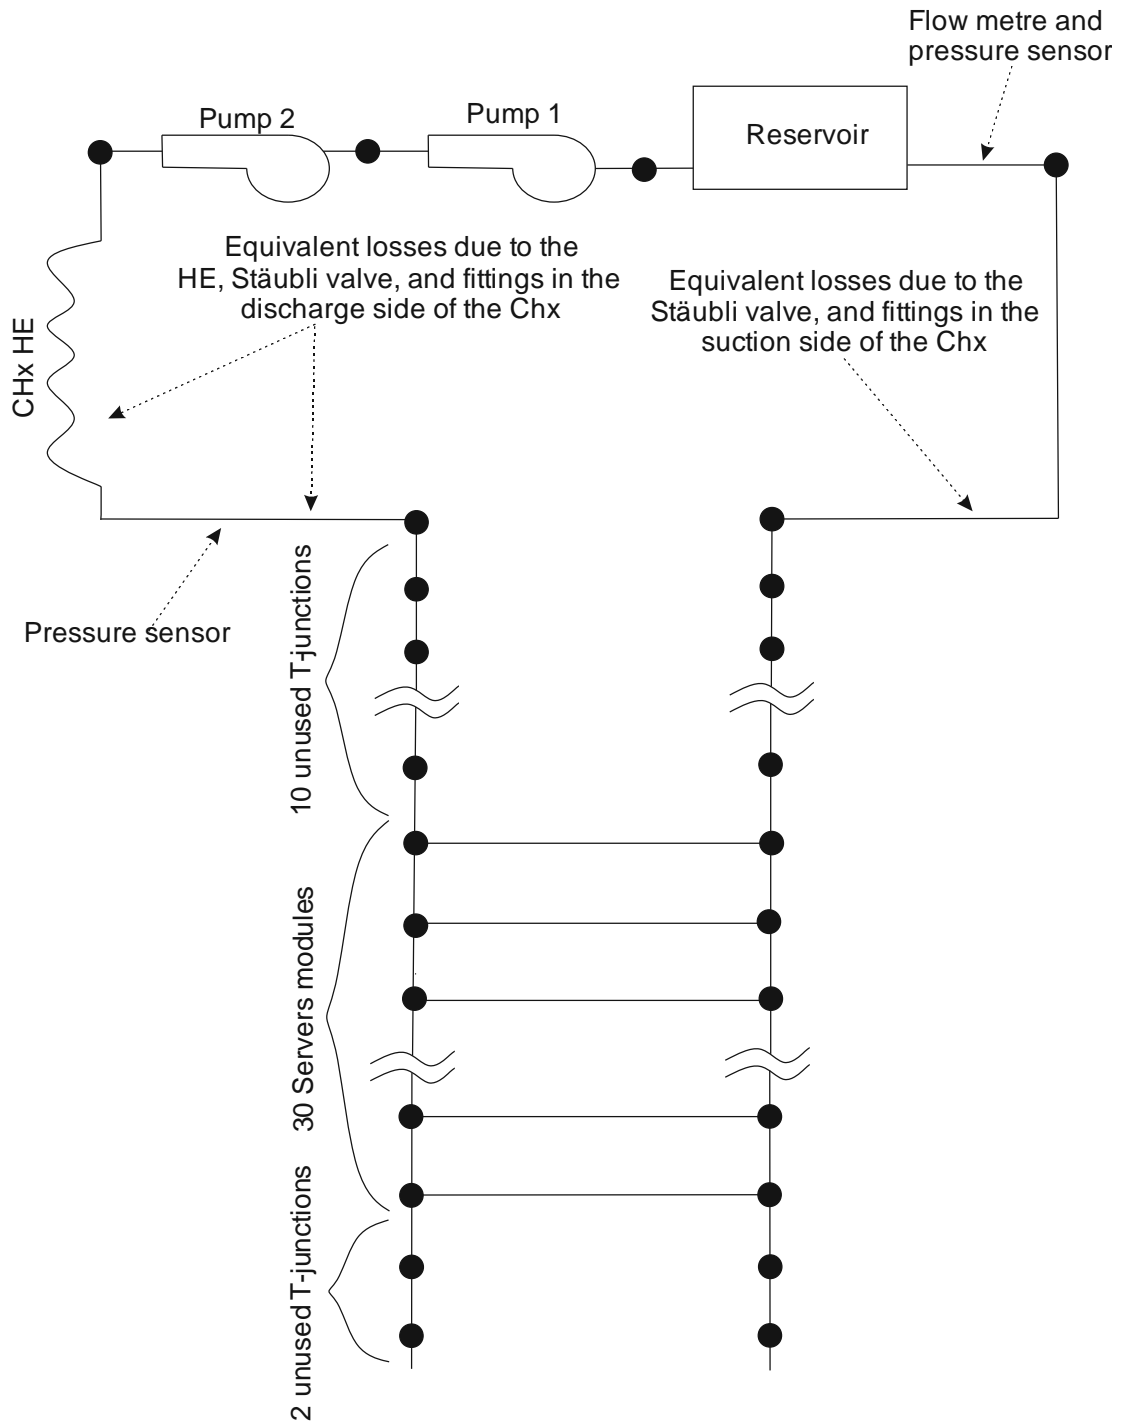
Figure 4.5 Secondary loop design for centralised pumping in EPANET software
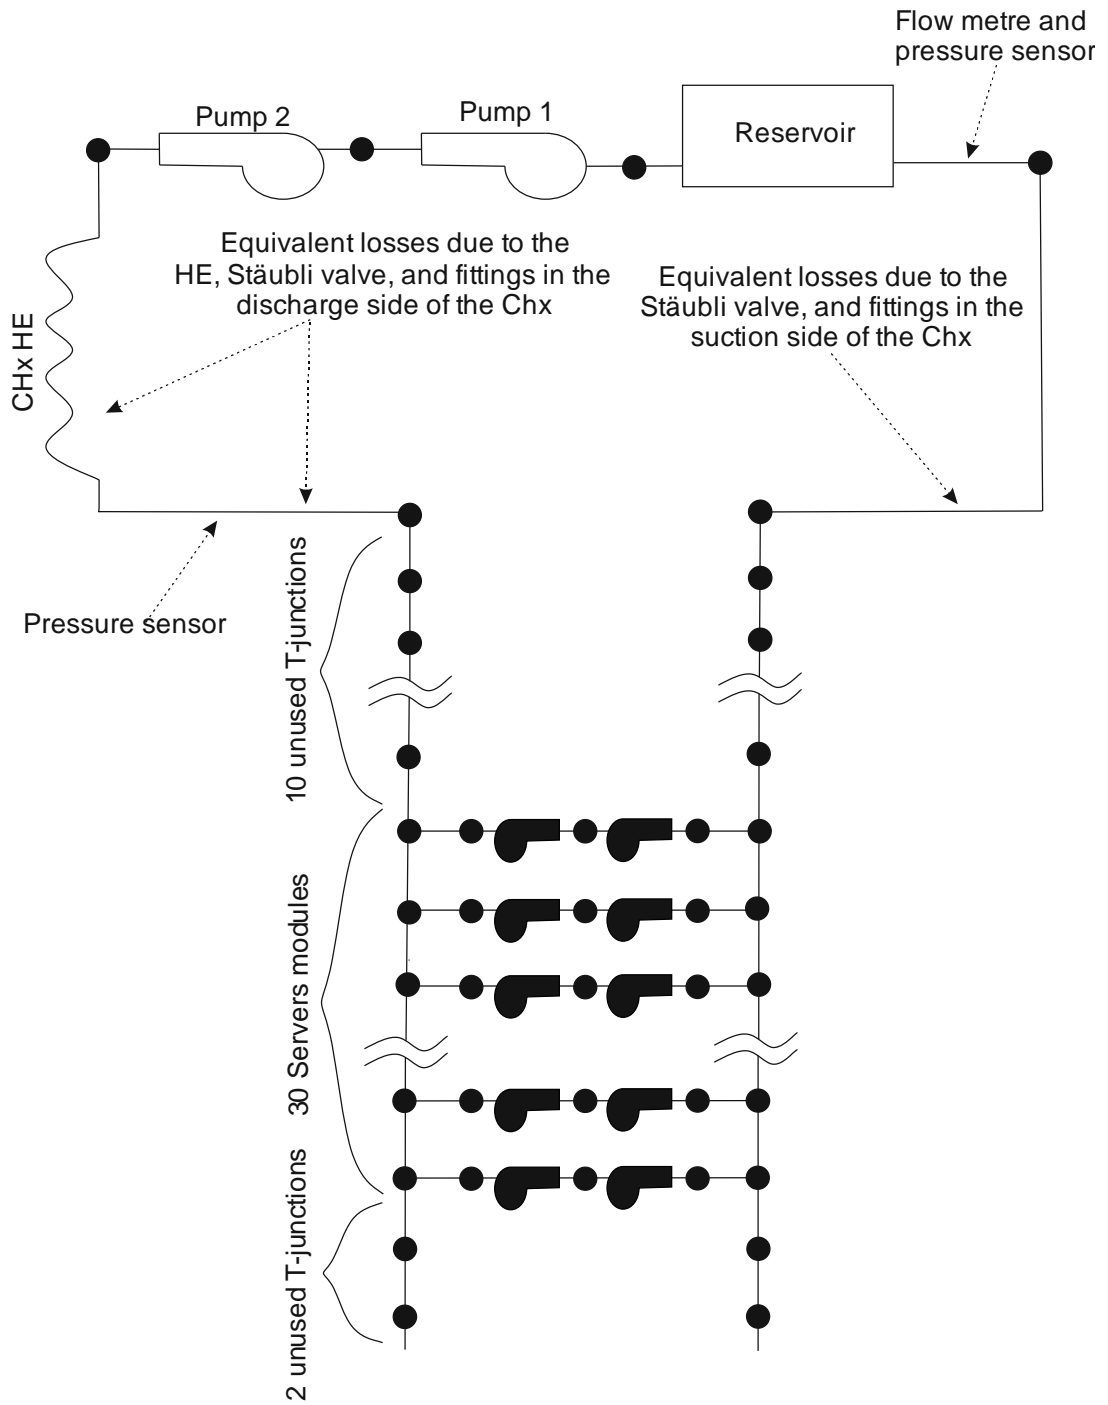
Figure 4.6 Secondary loop design for distributed pumping in EPANET software
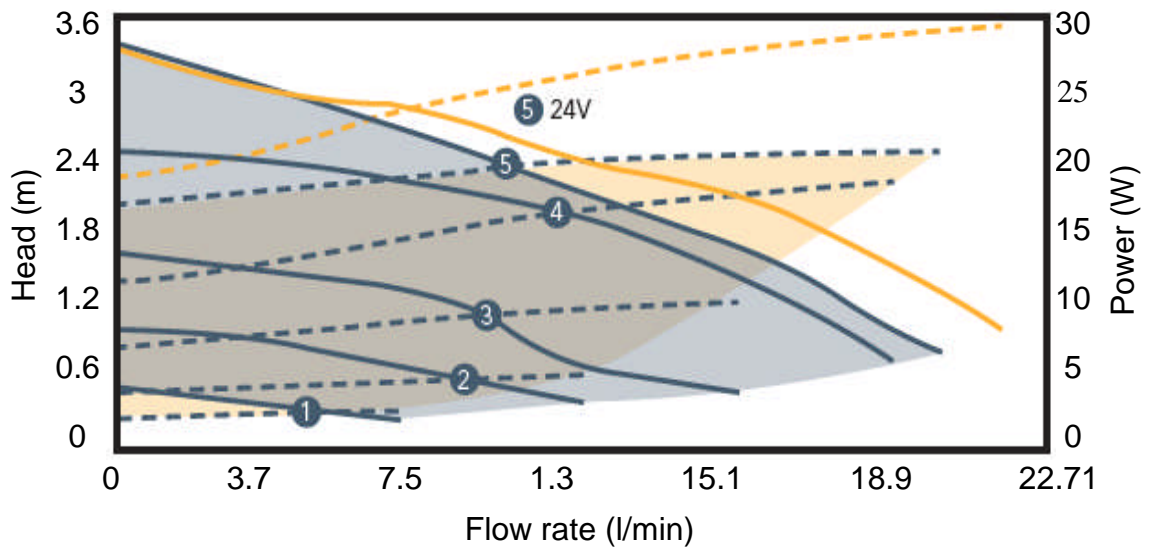
Figure 4.7 Characteristics curve of the central pumps; the blue shaded area represents the pump head against the flow rate at 12V while the yellow shaded area represents the power consumption. The solid and dashed yellow curves represent the pump head and the power consumption at 24V, respectively. The number on the curve refers to the pump speed setting (D5 series-Brass VARIO pump) [149].

4.11 Effect of rack inlet water temperature

The goal of testing the effect of different inlet water temperatures to the rack is to collect thermal, power, and performance data to validate the DCLC servers in a rack system for efficient cooling. To perform the study, temperature readings are collected at various measurement points from CPUs to the external heat rejection system and power consumptions of all the systems are measured using the same approach mentioned earlier in this chapter. All the remaining reading and data processing which are presented through this chapter are followed to calculate the different aspects and performance metrics of the rack operation. The inlet temperature to the rack is controlled using the controller in the AHU, which is set in auto mode of operation to supply the required temperature for different conditions of operation.

120-hours of experiments with three repeated trials were performed under different inlet water temperatures to the rack varying from $15℃$ to $40℃$ in increments of $5℃$. It should be noted that this temperature range of control is representing the inlet temperature of the primary loop to the rear door before passing through the CHx40 primary loop and not the inlet temperature to the servers. This is to highlight that the study is to investigate the effect of inlet temperature to the servers within the ASHRAE W4 liquid cooling recommendation range [52], ranging from $2℃$ to $45℃$. The rack was loaded computationally using the aforementioned scripts with idle, 25%, 50%, 75%, and 100% CPU stress to provide a complete set of data under different data centre loadings. Each computational stress level was run for 8 hours, and the inlet temperature point was changed every one hour starting in the third hour of the stress to allow the system to reach steady state. The data of the last 20 minutes of every hour was considered as the steady state data which was taken for further analysis.

Chapter 5 Operational performance of the data centre cooling system

5.1 Introduction

Bringing the liquid in direct contact with the heat generation sources in data centres with economisation could potentially yield significant improvement in data centre energy metrics. The proposed bespoke AHU and the cooling facility was designed to provide a way to quantify possible energy savings and provide a reliable means of cooling for HPC systems under a wide range of ambient conditions and IT loads.
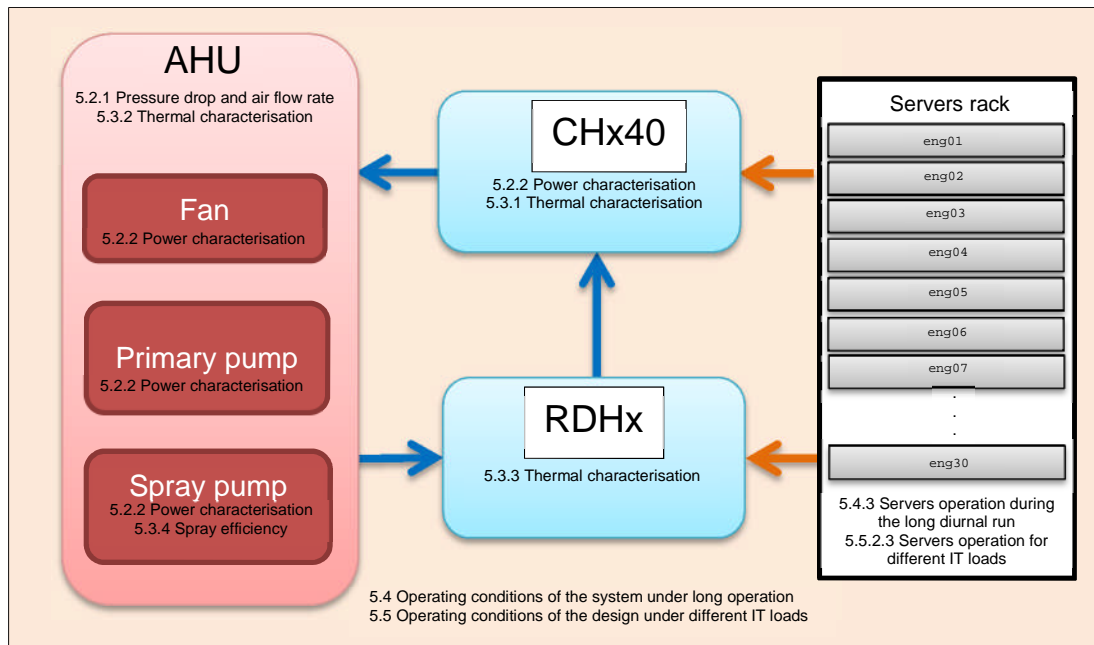
Figure 5.1 Block diagram showing the general structure of the sections of this chapter.

In this chapter, the results of several experimental tests on the designed data centre cooling facility are reported and analysed. The detailed analysis of the dynamic responses of each heat exchange component in the loop and the whole system are presented under varying IT loads and weather conditions. The thermal and power characteristics of each component in the system will be determined as a function of the primary flow rate, secondary flow rate, fan speed, and the spray flow rate. The study will also cover the results of long operation of the system under various ambient conditions and the

accompanied effect on power consumption and thermal characteristics. Finally, the different operational scenarios of the rack load will be considered to evaluate the performance of the designed cooling system under possible operation loads of a real data centre. Figure 5.1 shows the general structure of this chapter.

5.2 Impact of operation conditions on power consumption, flow conditions, and pressure drop in the designed facility

5.2.1 Air flow rate and pressure drop in the AHU

As mentioned in the previous chapter, the air flow rate is measured based on the pressure drop across the fan and the characteristic curve which was calibrated by the traversing method. Figure 5.2 shows the relationship between the fan speed in rpm and the volumetric air flow rate in m^3/min. The fitted data shows a polynomial relationship, as in equation 5.1, with a regression factor of 0.99.

$$Q_a = aF_s + bF_s^2 \qquad\qquad 5.1$$

Where, Q_a is the air flow rate in (m^3/min), F$_s$ is the fan speed in (rpm), and a and b are the constants of the equation which are respectively; 21.39 and 0.012.

The pressure drop across the AHU HE is shown in figure 5.3 for the dry and wet cases as the fan speed is varied. The pressure drop increases 33% when increasing the fan speed from 250 rpm to 1000 rpm in the dry operation. However, the addition of the spray shows a significant increase in the pressure losses in all the tested cases and a maximum increase of about 145% was recorded at the higher fan speed and spray flow rate. Figure 5.4 shows the relationship between the pressure drop, fan speed and the spray flow rate. The data has been fitted to a quadratic model with a regression factor of 0.98, using the Design of Expert software (DoE V10) [144], as in equation 5.2. The comparison between the predicted results of the model (equation 5.2) and the actual data obtained from the experiments is shown in figure 5.5.

$$\Delta p = \frac{1}{(a + bF_s + cQ_s + dF_s Q_s + eF_s^2 + gQ_s^2)^2}$$

5.2

Where Δp and Q_s are the pressure drop across the AHU HE and the spray flow rate in (l/min), respectively.

Table 5.1 constants of equation 5.2

a	b	c	d	e	g
0.29431	-9.4094 $\times 10^{-5}$	-4.497 $\times 10^{-4}$	4.545 $\times 10^{-7}$	1.792 $\times 10^{-8}$	-8.3203 $\times 10^{-5}$
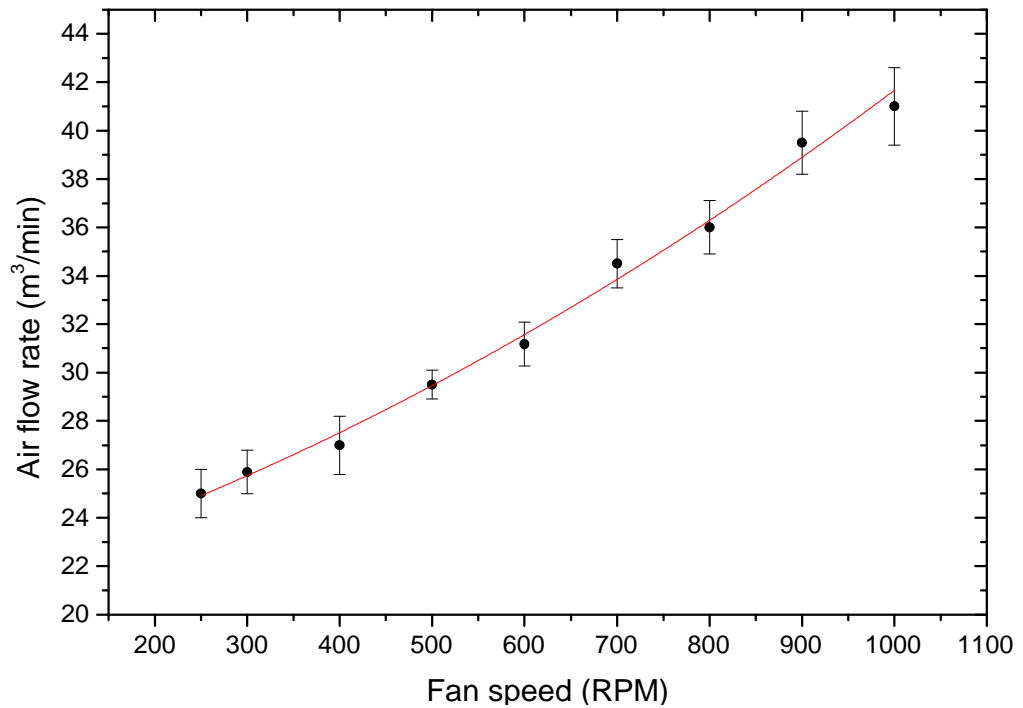
Figure 5.2 The air flow rate vs the fan speed of the AHU.
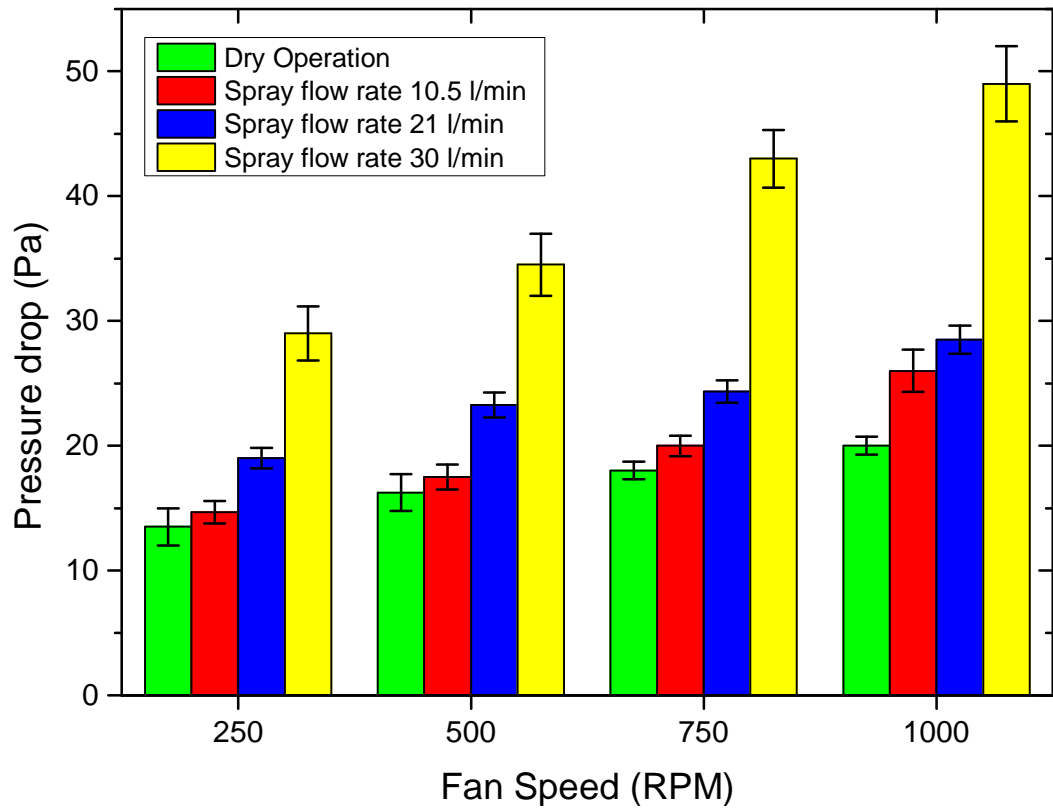
Figure 5.3 Pressure drop across the AHU heat exchanger as a function of the fan speed for different spray water flow rates.
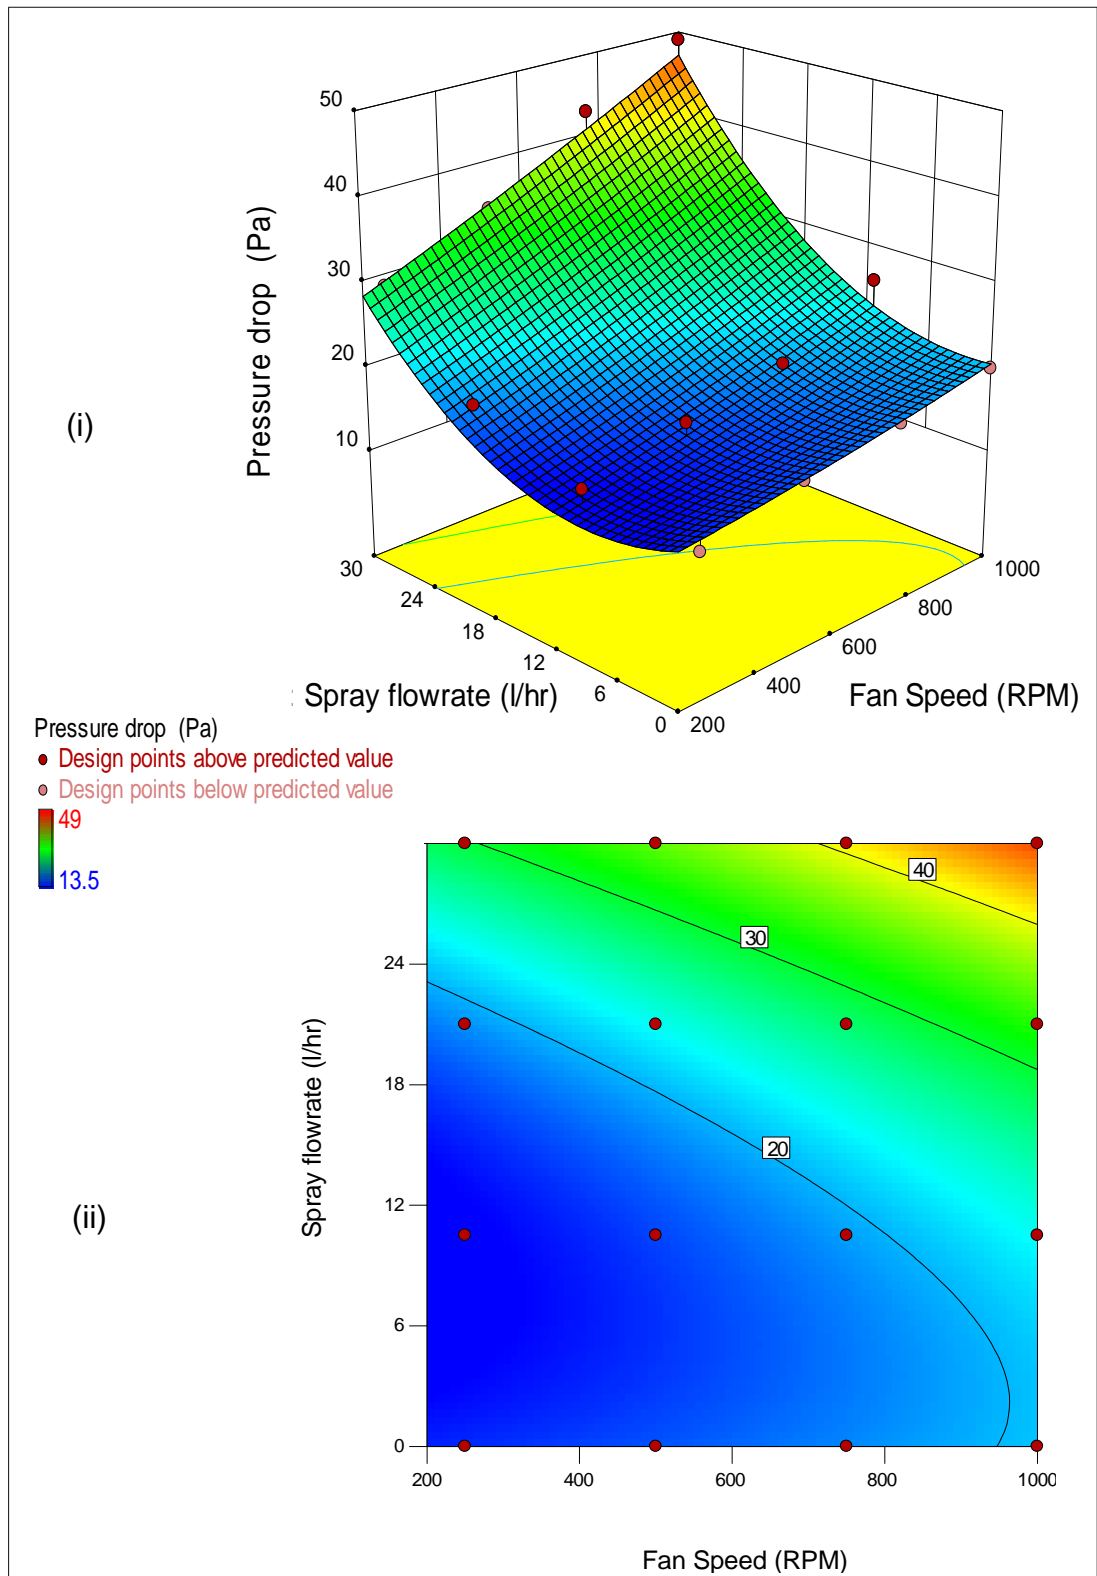
Figure 5.4 Approached pressure drop surface for the AHU heat exchanger as a function of fan speed and spray flow rate; (i) 3D graph showing the surface, and (ii) 2D graph showing the measurement points.
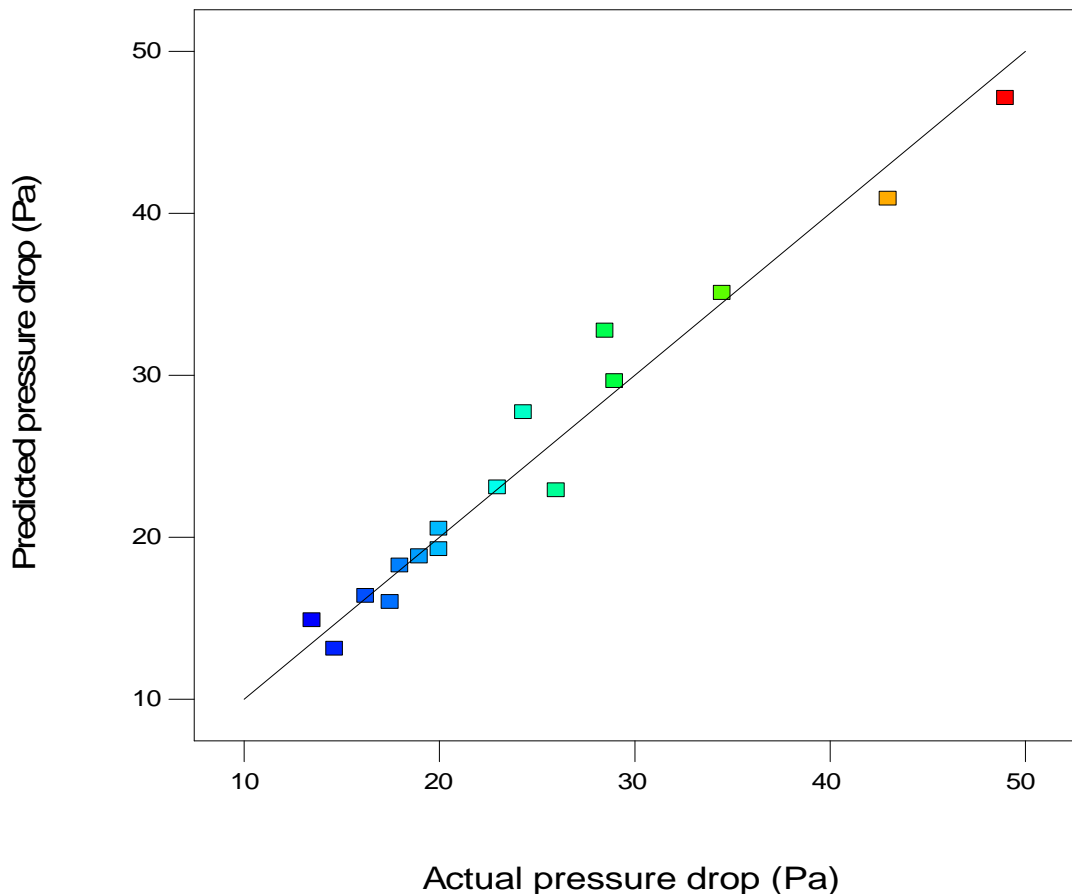
Figure 5.5 The actual vales of the pressure drop from the experiments against the predicated values using equation 5.2.

5.2.2 Power consumption analysis of the designed cooling system

Evaluation of the power consumption is necessary to calculate the effectiveness of the system and to have an insight of how much each component consumes in the cooling system. The pressure drop in each loop is also associated with an increase in the power consumption of the components and the data centre as a whole in general. Figure 5.6 shows the power consumption by the CHx40 unit as the secondary flow rate is varied. The CHx40 consumes about 8W even when the pumps are off. It can be seen in figure 5.6, the difference in the power consumption between the central and distributed pumping. For the same flow rate, the power consumption in the CHx40 unit is higher with central pumping rather than with distributed pumping. Unsurprisingly, this implies that the pPUE for central pumping is

higher than the distributed pumping for the same flow rate because the part of the power consumed by the cooling side will be transferred to the IT side according to the pPUE relationship.

The power consumption of the primary loop pump and the spray pump is shown in figures 5.7 and 5.8, respectively. The power consumption of the primary loop pump varies from $\sim 5W$ at $3\,l/min$ to $\sim 255W$ at the maximum flow rate of $19\,l/min$. However, the spray pump has a higher power consumption where it consumes $\sim 175W$ at the minimum flow rate of $2.5\,l/hr$ up to $\sim 300W$ at $30\,l/hr$. The reason of high power consumption of the spray unit is that the pumps need to overcome the high pressure drop associated with the nozzles and the solenoid valves. As a result, the addition of the spray increases the power consumption, water usage, and pressure drop in the data centre cooling system. Consequently, using spray within the loop must be weighed against the thermal benefit which is obtained from the evaporative cooling.

Finally, the power consumed by the AHU fan over the speed range of 100 to 1000 rpm is shown in figure 5.9. The power consumption of the fan is very low at 100 rpm of $5W$ and increases rapidly with increasing the fan speed to reach $960W$ at the maximum speed limit. The is due to the high pressure drop associated with increasing the fan speed which is a result of the HE, ducts, damper and the building exhaust duct resistances. The power consumed by the fan is comparable with the power consumed by the primary loop pump at speeds up to about 600 rpm. However, beyond that limit, the power consumption by the fan increases rapidly which makes the operation undesirable with fan speeds greater than 600 rpm. However, under some conditions it is required to use more power on the cooling to provide the required safe temperature to the data centre thus requiring a balance between higher fan speeds and activating the spray. The associated losses of using the spray is not only the power consumption of the pump and the solenoid valves but also there is an associated water loss which adds requirement for water treatment, water consumption, and the potential risks of using the spray which are associated with the health and safety. Thus, the spray was designed

to be activated only when the fan reaches 800 rpm to reduce the increase in the fan power consumption and to increase the thermal capacity of the AHU HE when the fan cannot on its own manage the cooling load. This implies that using the spray under some cases could have a significance on improving the pPUE metric of the data centre but at the same time it has a negative impact on the water usage effectiveness and the total operation cost of the data centre.
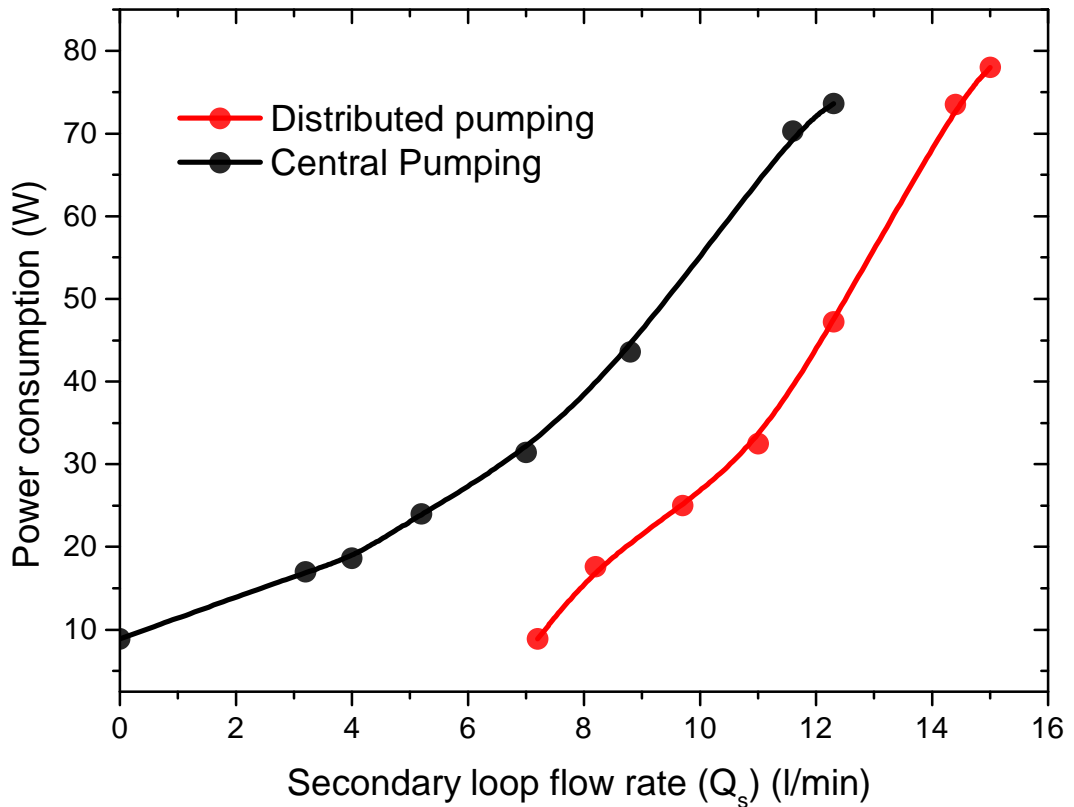
Figure 5.6 CHx40 Power consumption as a function of the secondary loop flow rate for the central and distributed pumping configurations.
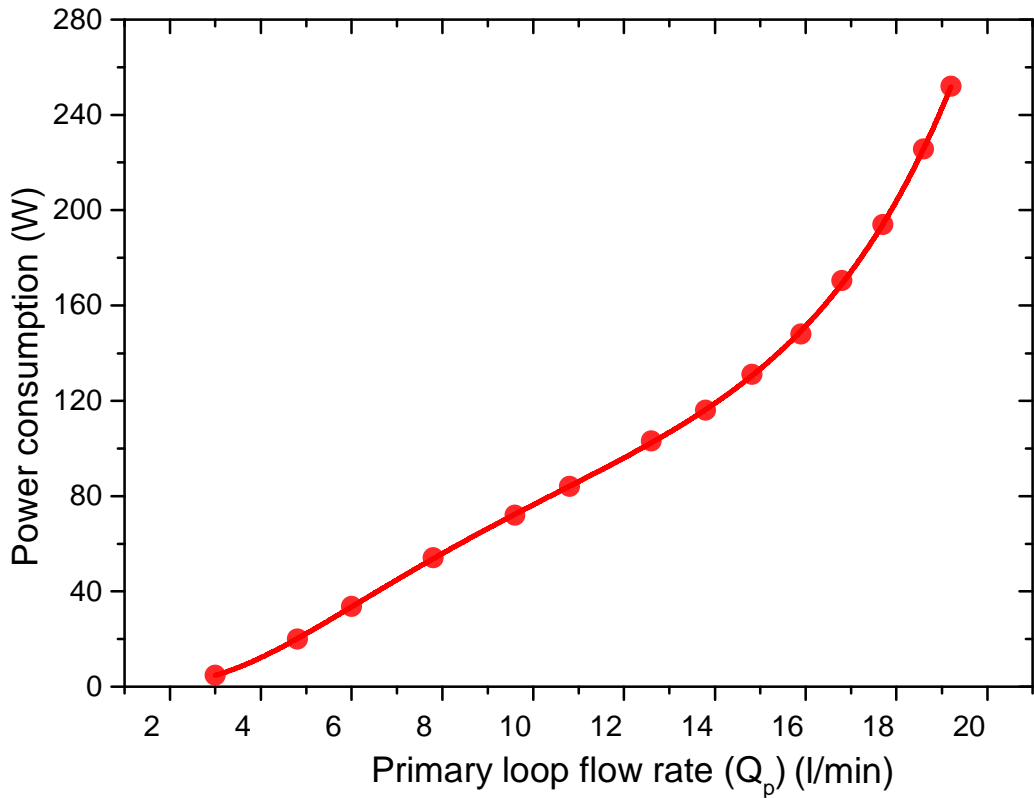
Figure 5.7 Primary loop pump power consumption as a function of the primary loop flow rate.
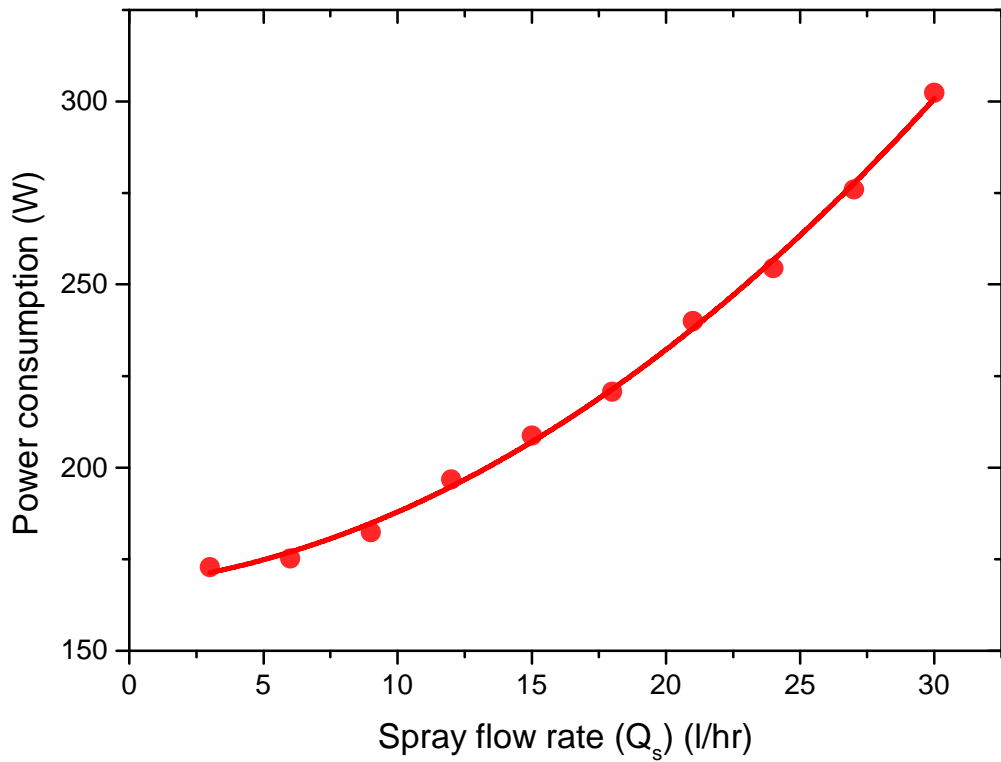
Figure 5.8 Spray pump power consumption as a function of the spray flow rate.
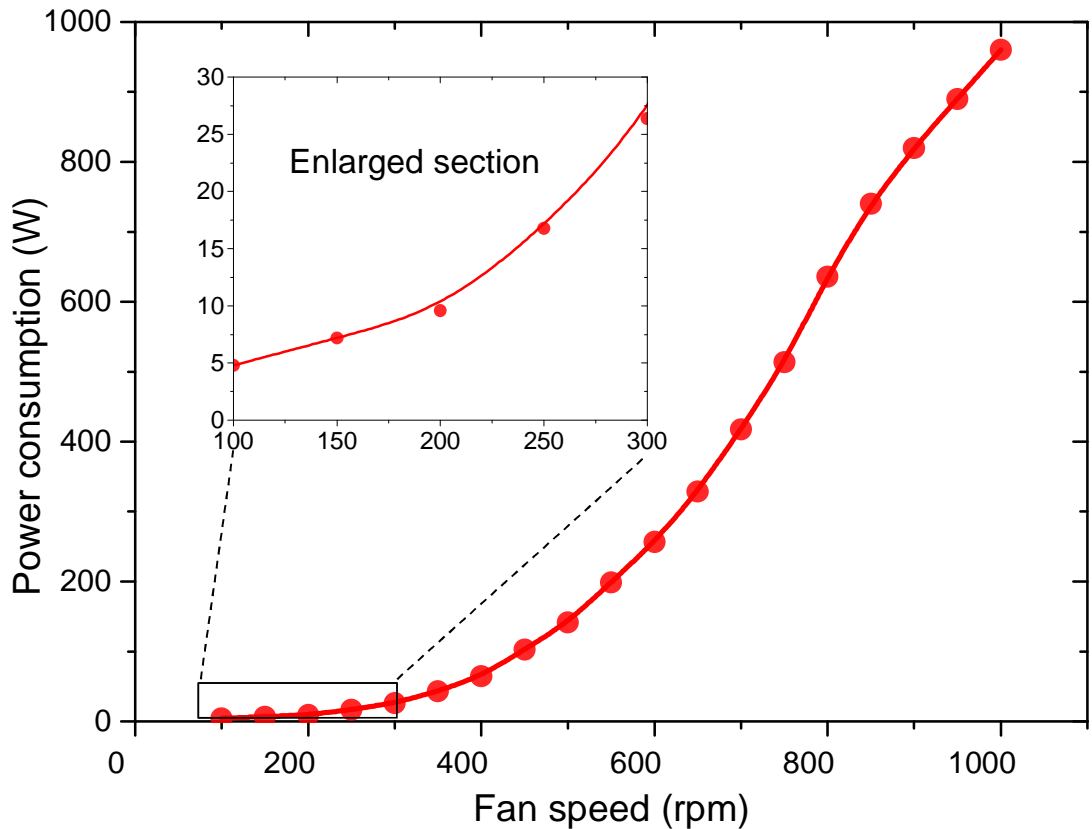
Figure 5.9 AHU fan power consumption as a function of fan speed.

5.3 Characterisation of the cooling system heat exchange points

5.3.1 CHx40 heat exchanger

As mentioned earlier, the CHx40 heat exchanger enables heat exchange between the secondary loop with the servers and the primary loop with the AHU to reject the heat to the environment. The primary loop coolant is a mixture of water and propylene glycol to avoid freeze damage in the AHU HE in winter. The secondary loop coolant is a mixture of water and ethylene glycol to increase corrosion resistance of the components. The impact of flow variation in the secondary and primary loops on the thermal resistance of the CHx40 is shown in figure 5.10. The results are best fitted by a surface quadratic model in equation 5.3 with a regression factor of 0.98 using the Design of Expert software (DoE V10) [144]. The fitted R_{CHx} values against the measured values from experiment is shown in figure 5.11.

$$R_{CHx} = \frac{1}{\left(a + bQ_p + cQ_s + dQ_pQ_s + eQ_p^2 + fQ_s^2\right)^2}$$

5.3

Where the constants of the equation are given in table 5.2

Table 5.2 Constants of equation 5.3

a	b	c	d	e	f
0.167	0.0241	0.0308	1.175×10^{-3}	-1.073×10^{-3}	-1.0889×10^{-3}

It can be seen from figure 5.10 that the higher thermal resistance of the CHx40 HE is at lower flow rates of the primary and secondary loops. The thermal resistance improves as the primary and secondary flow rates increase. The decrease or increase of the thermal resistance depends mainly on the conductance of the heat exchanger as it can also be expressed as a function of the HE effectiveness and the heat capacities of the secondary and primary loops as in equation 5.4. Hence, if the HE conductance term is high, then the thermal resistance will be improved with increasing the secondary loop flow rate which is the case in the CHx40. However, the thermal resistance may increase with increasing the secondary loop flow rate depending on the thermal conductance of the HE.

$$R_{CHx} = \frac{1}{\varepsilon.C_s} - \frac{1}{C_p}$$

5.4

Where ε, C_s and C_p are the effectiveness, heat capacity of the secondary loop, and the heat capacity of the primary loop, respectively.
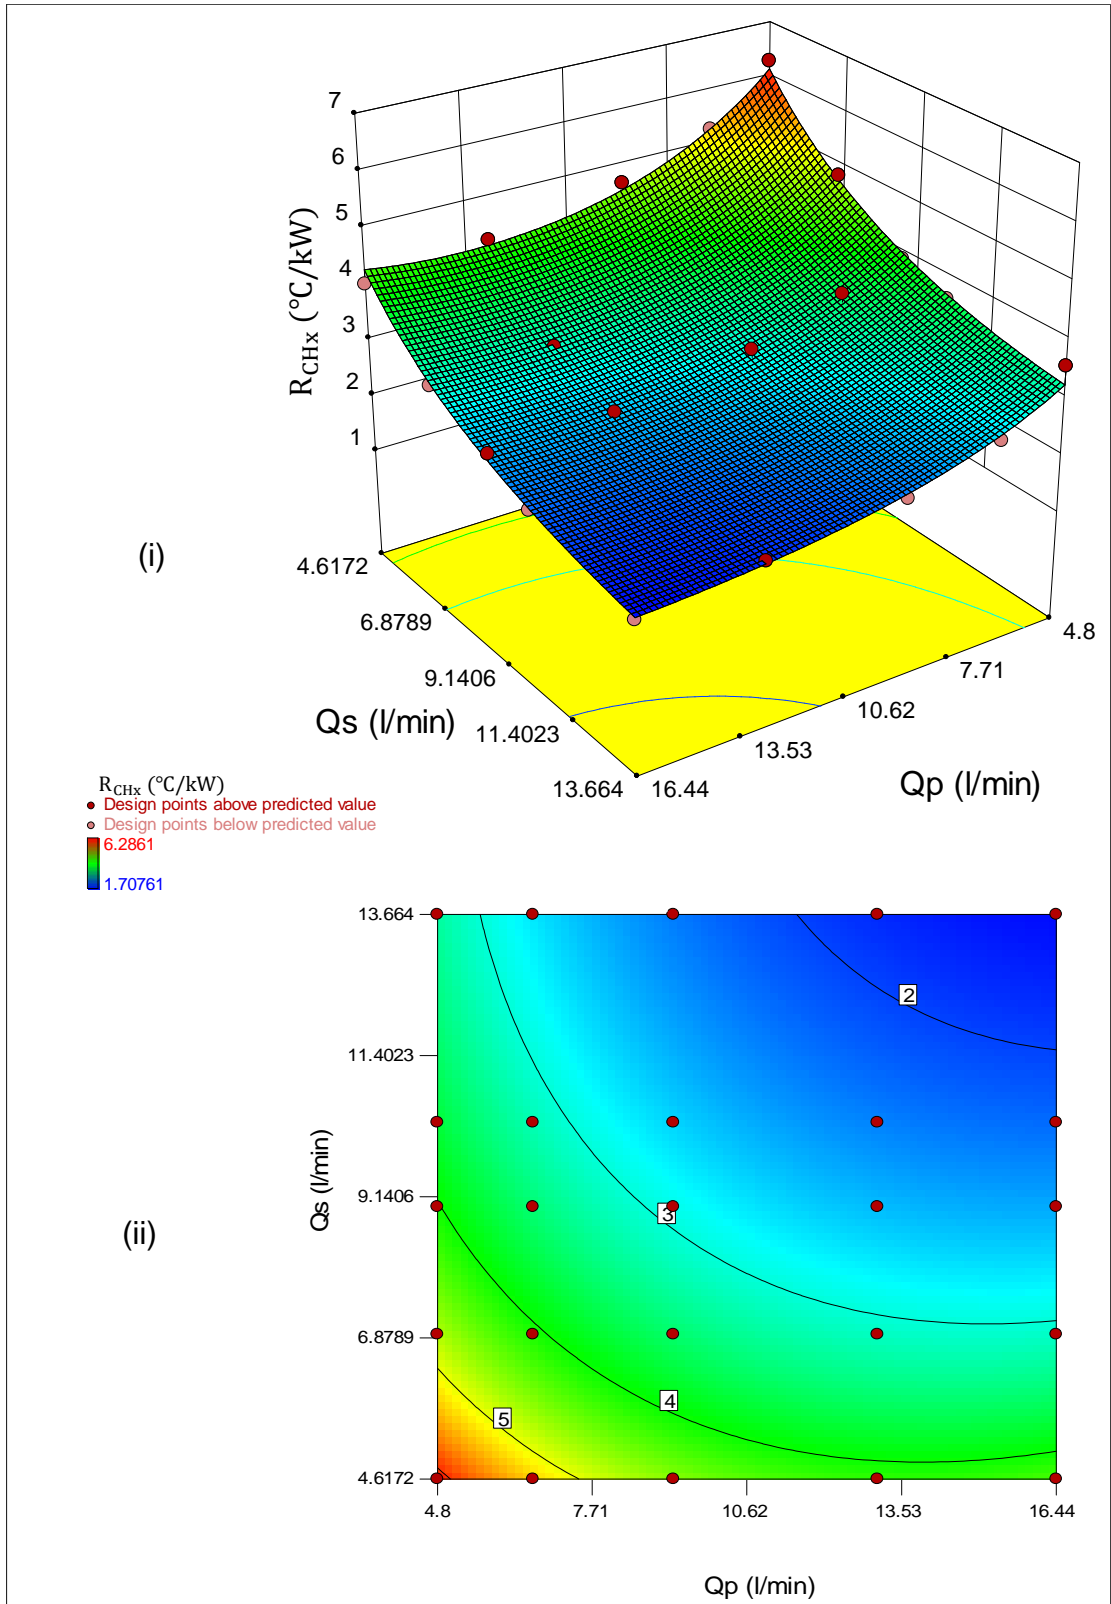
Figure 5.10 Thermal resistance of the CHx40 HE as a function of primary (external) (Q_P) and secondary (internal) (Q_s) flow rates; (i) Shows the predicted 3D graph and (ii) shows the predicted and the actual points on 2D graph.
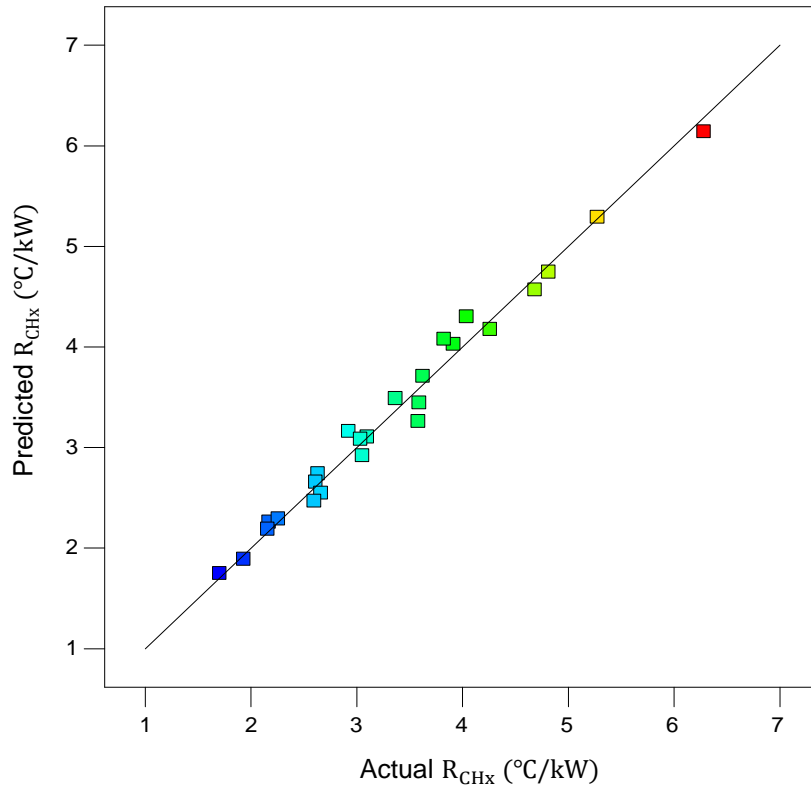
Figure 5.11 The actual values of the R_CHx from the experiments against the predicated values using equation 5.3.

5.3.2 Thermal resistance of the AHU

The AHU is used to exchange the heat in the primary coolant loop, which is generated by the IT, to the environment. Like the CHx40, the thermal performance of the AHU depends on the flow rates of the air and the coolant in the primary loop as well as activating the spray in the AHU.

Figure 5.12 shows the impact of the primary coolant flow rate and the fan speed in rpm on the approach thermal resistance of the AHU HE in the dry case. The best fit surface curve is determined, as shown in equation 5.5, based on the surface quadratic model. The regression factor of the model is shown to be 0.99 and calculated values against the predicted values curve are shown in figure 5.13.

$$R_{AHU\_dry} = a + bQ_p + cF_s + dQ_pF_s + eQ_p^2 + fF_s^2 \qquad 5.5$$

The constants are shown in table 5.3

Table 5.3 Constants of equation 5.5

A	b	c	d	e	f
5.59	−0.211	-7.393×10^{-4}	-1.705×10^{-5}	2.035×10^{-3}	2.901×10^{-7}

The thermal resistance is higher at the lower coolant flow rate and fan speed, and strongly reduces with increasing primary loop flow rate. The thermal resistance also decreases with increasing fan speed. However, the effect of increasing the primary loop flow rate on the thermal benefits outweigh the effect of increasing the fan speed (air flow rate). Similar to the CHx40 HE, the thermal resistance of the AHU HE is also a function of the effectiveness and thermal capacities of the primary coolant loop and the air flow rate as shown in equation 5.6.

$$R_{AHU\_dry} = \frac{1}{\varepsilon.C_a} - \frac{1}{C_p}$$

5.6

C_a is the heat capacity of the air loop in the AHU.

In the current case, increasing the primary loop flow rate results relatively in a higher reduction in the term $\left(\frac{1}{\varepsilon.C_a}\right)$ and compared to the reduction in the term $\left(\frac{1}{C_p}\right)$ due to the high thermal conductance of the HE exchanger surface. However, this is not always the case as David et al [96] found that there is a deterioration in the thermal resistance of a dry cooler with increasing flow rate in the liquid side.
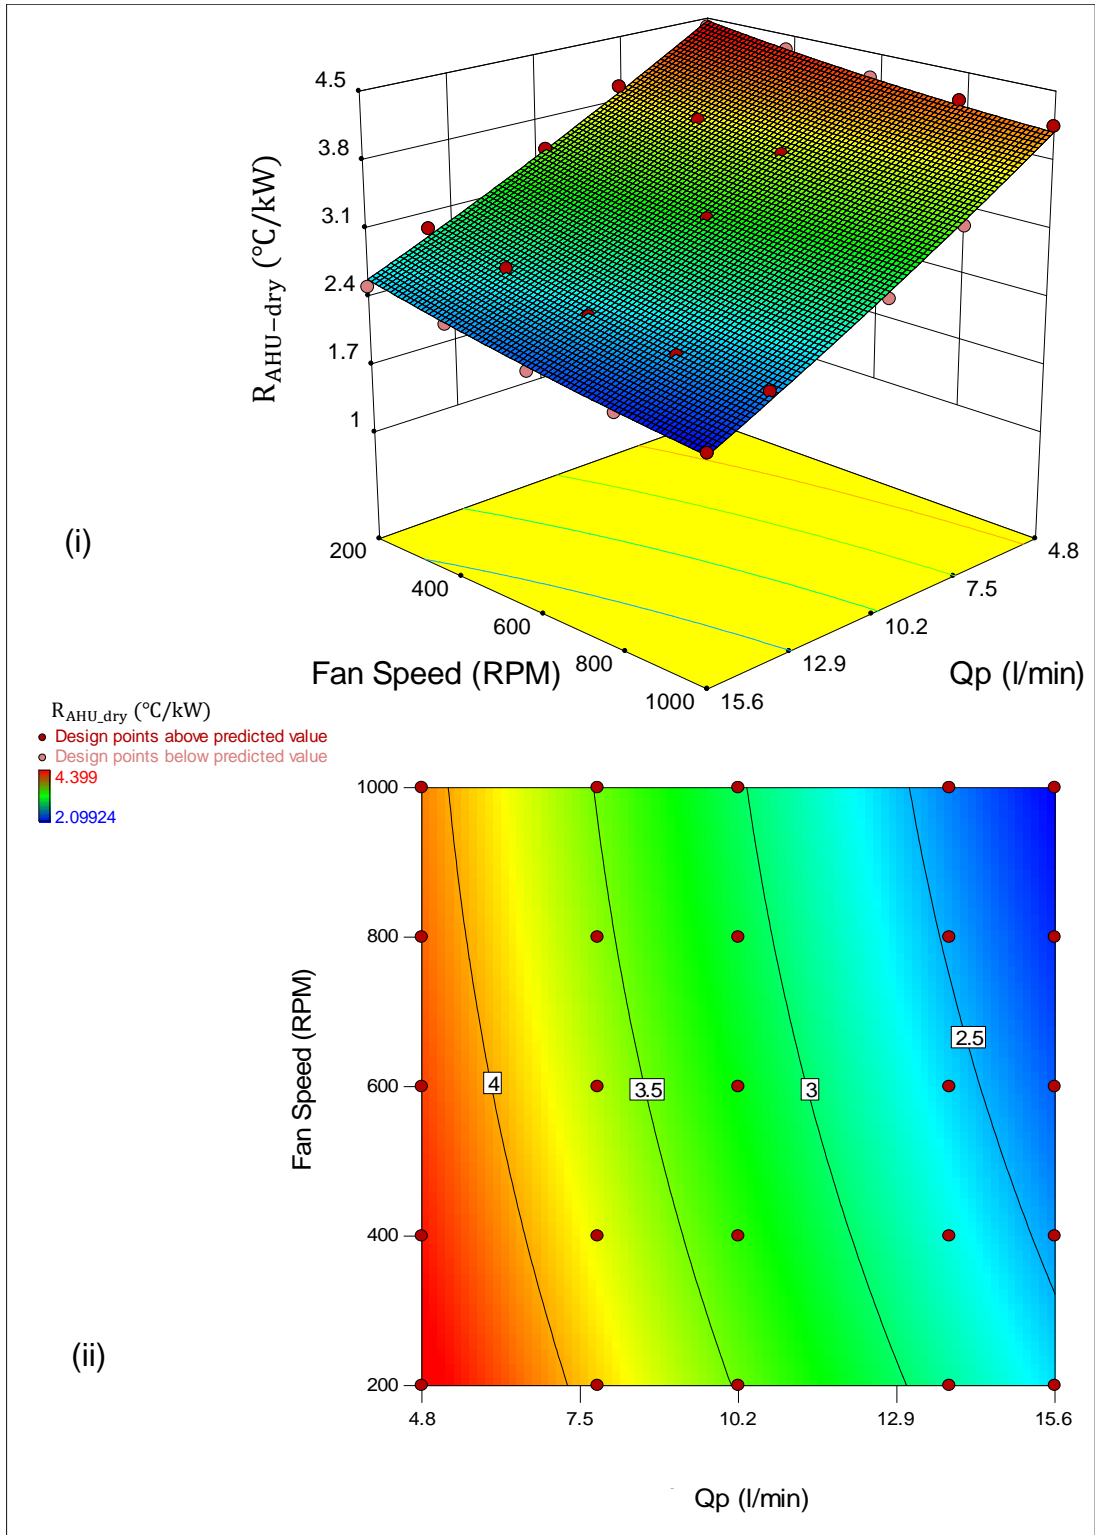
Figure 5.12 Thermal resistance of the AHU HE as a function of the primary loop flow rate and the AHU fan speed in the dry case; (i) Shows the predicted 3D graph and (ii) shows the predicted and the actual points on 2D graph.
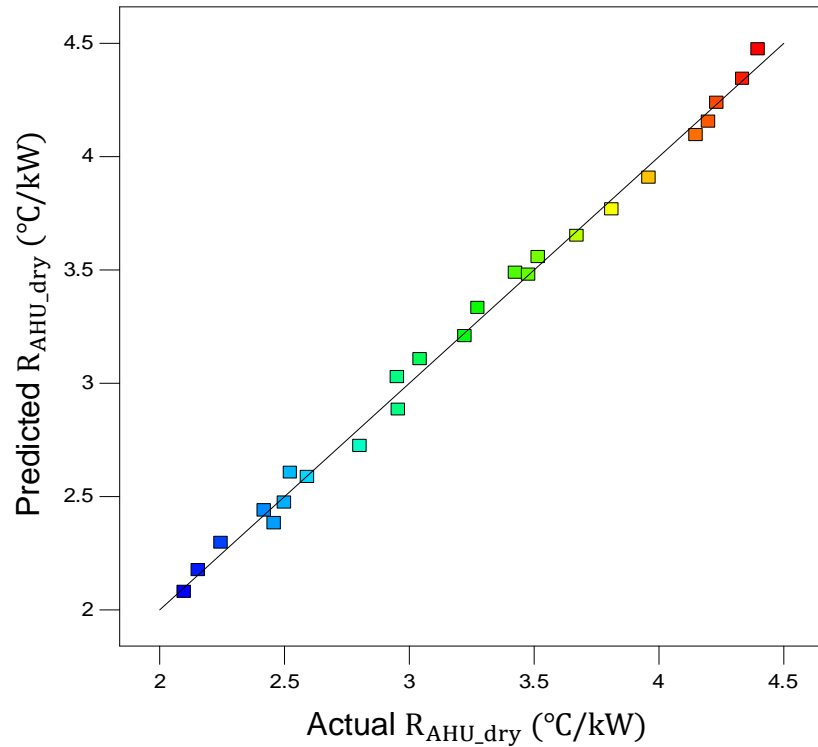
Figure 5.13 Fitted R_{AHU_dry} from equation 5.5 vs the actual values calculated from the experiment in the dry case.

Spray operation of the AHU was also tested under different fan speeds and primary coolant flow rate. In these test conditions, the exchange from the primary coolant loop to the environment is by two means; sensible and latent heat transfers due to the spray droplet evaporation. Due to the much higher heat transfer when the spray is activated, the thermal resistance is much lower than the dry case. Activation of the spray reduces the thermal resistance by about 50% at the lower fan speed and primary loop flow rate. Figure 5.14 shows the impact of the spray on the thermal resistance of the AHU HE. The effect of the fan speed on the thermal resistance is lower than the dry case and it is almost insignificant at the lowest primary water flow rate. The best fit of the surface is shown by the surface quadratic model in equation 5.7. The model shows a regression factor of 0.98 and predicted values are shown against the actual data in figure 5.15.

$$R_{AHU\_wet} = \frac{1}{\left(a + bQ_p + cF_s + dQ_pF_s + eQ_p^2 + fF_s^2\right)^2}$$ 5.7

Table 5.4 Constants of equation 5.7

a	b	c	d	e	f
2.397	−0.159	-3.252×10^{-5}	-2.087×10^{-6}	5.032×10^{-3}	-1.145×10^{-8}
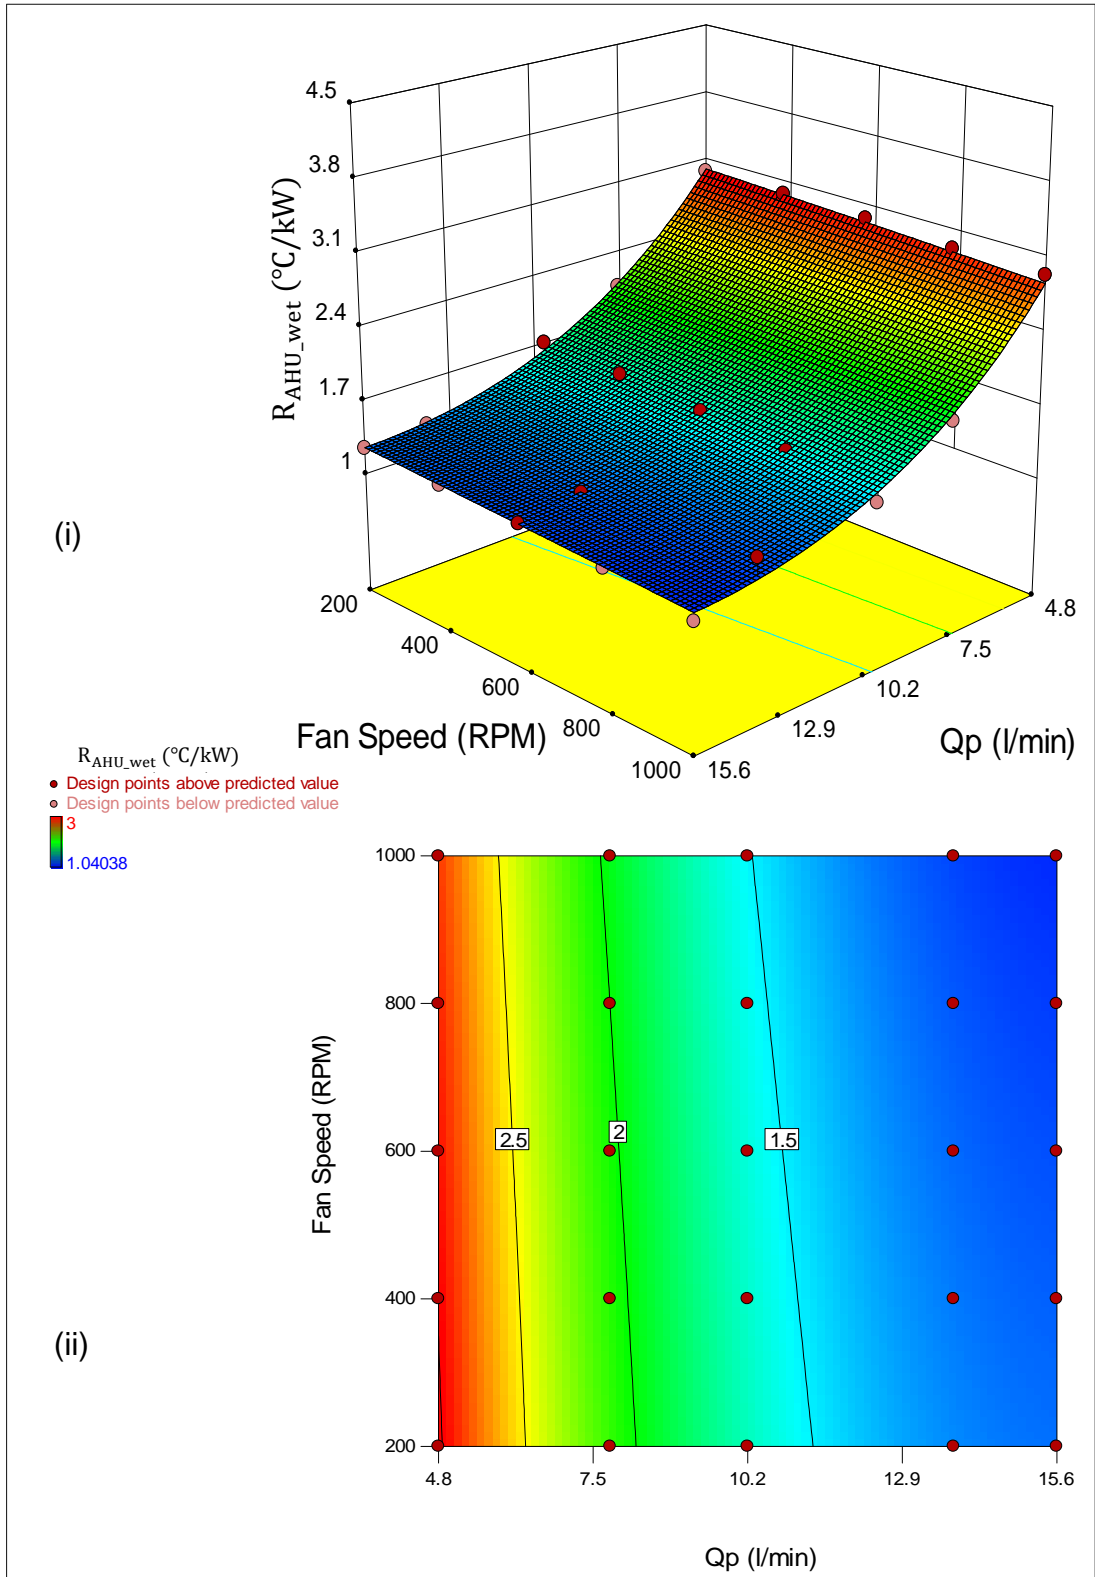
Figure 5.14 Thermal resistance of the AHU HE as a function of the primary loop flow rate and the AHU fan speed in the spray case; (i) Shows the predicted 3D graph and (ii) shows the predicted and the actual points on 2D graph.

147
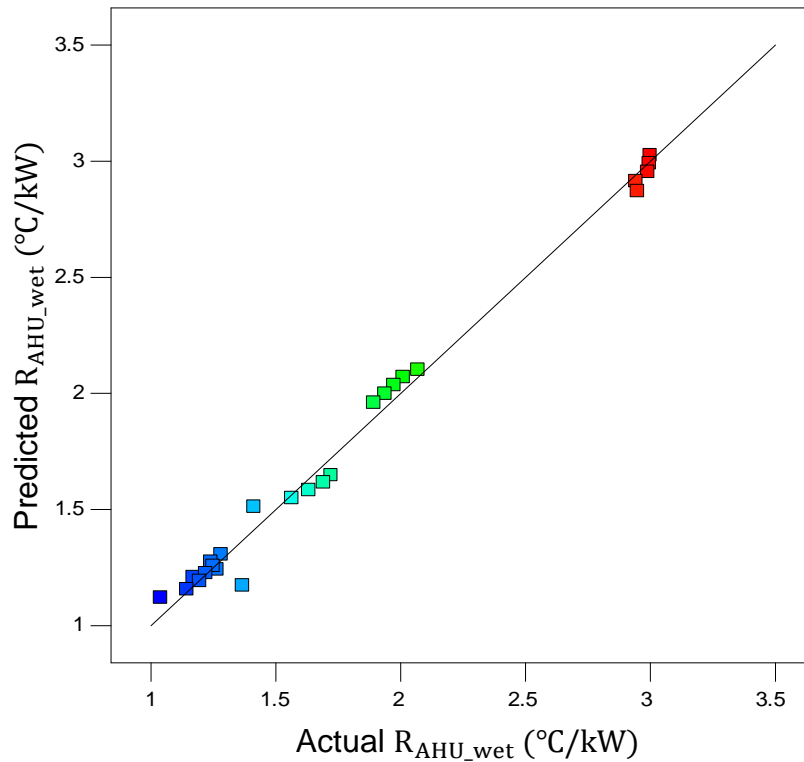
Figure 5.15 Calculated R_{AHU_wet} from equation 5.7 vs the actual values calculated from the experiment in the spray case.

5.3.3 Thermal operation of the RDHx

The rear door HE (RDHx) is used to exchange the heat generated by the RAM and the other auxiliary components of the servers with the process coolant loop. As mentioned in the previous chapters, the RDHx exchanges the heat in a passive way using the small fans in the servers to pass hot air through the (wavy) RDHx's fins. The servers' fan speeds are driven by a complex algorithm depending on the servers' temperature. Hence, it is difficult to measure and control the air flow rate through the RDHx. As a result, the effect of primary coolant flow rate on the thermal resistance of the RDHx was only investigated. Figure 5.16 shows the thermal resistance of the RDHx as a function of the primary loop flow rate. It can be seen that the thermal resistance of the HE decreases with increasing the flow rate and can be represented by a polynomial law as shown in equation 5.8 with a regression factor of 0.99. The results also show the influence of the high conductance of the RDHx which helps to decrease the thermal resistance with increasing primary water flow rate.

148

$$R_{RDHx} = aQ_p + bQ_p^2 + c \qquad\qquad 5.8$$

Table 5.5 Constants of equation 5.8

a	b	c
0.05864	−0.02651	11.436
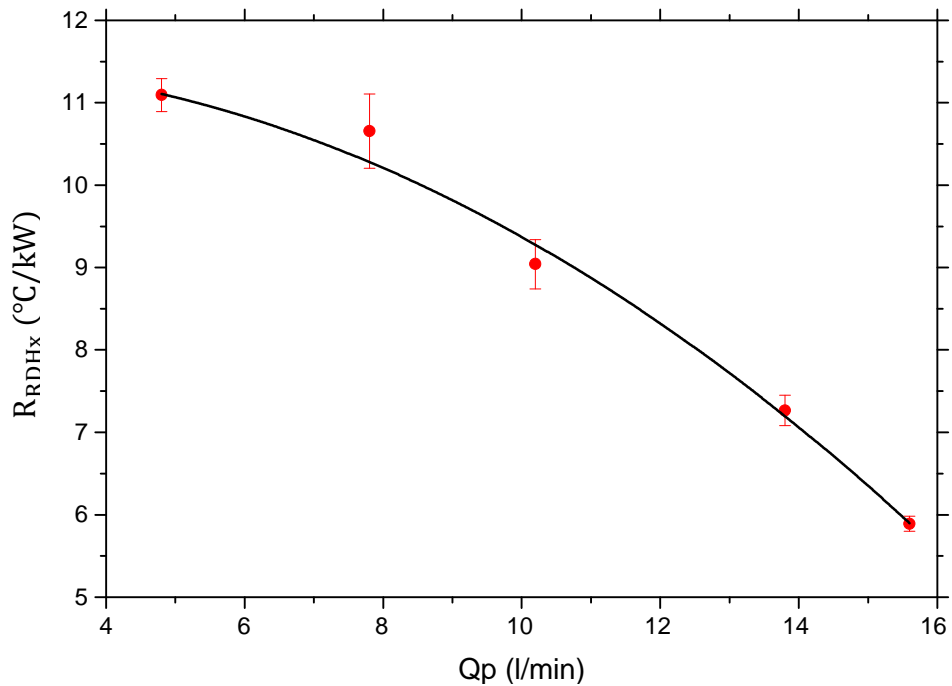
Figure 5.16 Thermal resistance of the rear door heat exchanger as a function of the primary loop flow rate.

5.3.4 Spray efficiency

The water sprayed over the heat exchanger reduces the thermal resistance. However, only part of the sprayed water evaporates and the other part is collected by the AHU sump. As a result, it is crucial to characterise the spray efficiency (SE) of the system to avoid over spraying and wasting water which increases the operational cost of the data centre and reduces the effectiveness by increasing the power consumption in the cooling side of the effectiveness equation. The spray efficiency depends on the primary loop flow rate and the fan speed. The higher fan speed and coolant flow rate, the higher rate of evaporation. Figure 5.17 shows the spray efficiency as a function of the AHU fan speed and coolant flow rate. The best fit surface is determined

and given by equation 5.9, with a regression factor of 0.98. The actual and predicted values by equation 5.9 are shown in figure 5.18.

$$SE = a + bQ_p + cF_s + dQ_pF_s + eQ_p^2 + fF_s^2 + gF_sQ_p^2 + hQ_pF_s^2 + iQ_p^3 + jF_s^3 \quad 5.9$$

The equation constants are given in table 5.6

Table 5.6 Constants of equation 5.9

a	b	c	d	e	f	g	h	i	j
-0.276	0.101	3.14×10^{-4}	-3.5×10^{-5}	-0.007	-1.17×10^{-7}	1.455×10^{-6}	4.326×10^{-10}	1.93×10^{-4}	7.193×10^{-11}

The results show that the maximum amount of water that can be effectively and potentially used to boost the heat exchanger capacity using the latent heat of evaporation is about 25% of the total sprayed water. This was at maximum flow rates of coolant and air. As a result, it is necessary to manage the use of the spray effectively when it is needed only in the cases when the ambient temperature is high or the AHU is unable to provide the required target temperature.
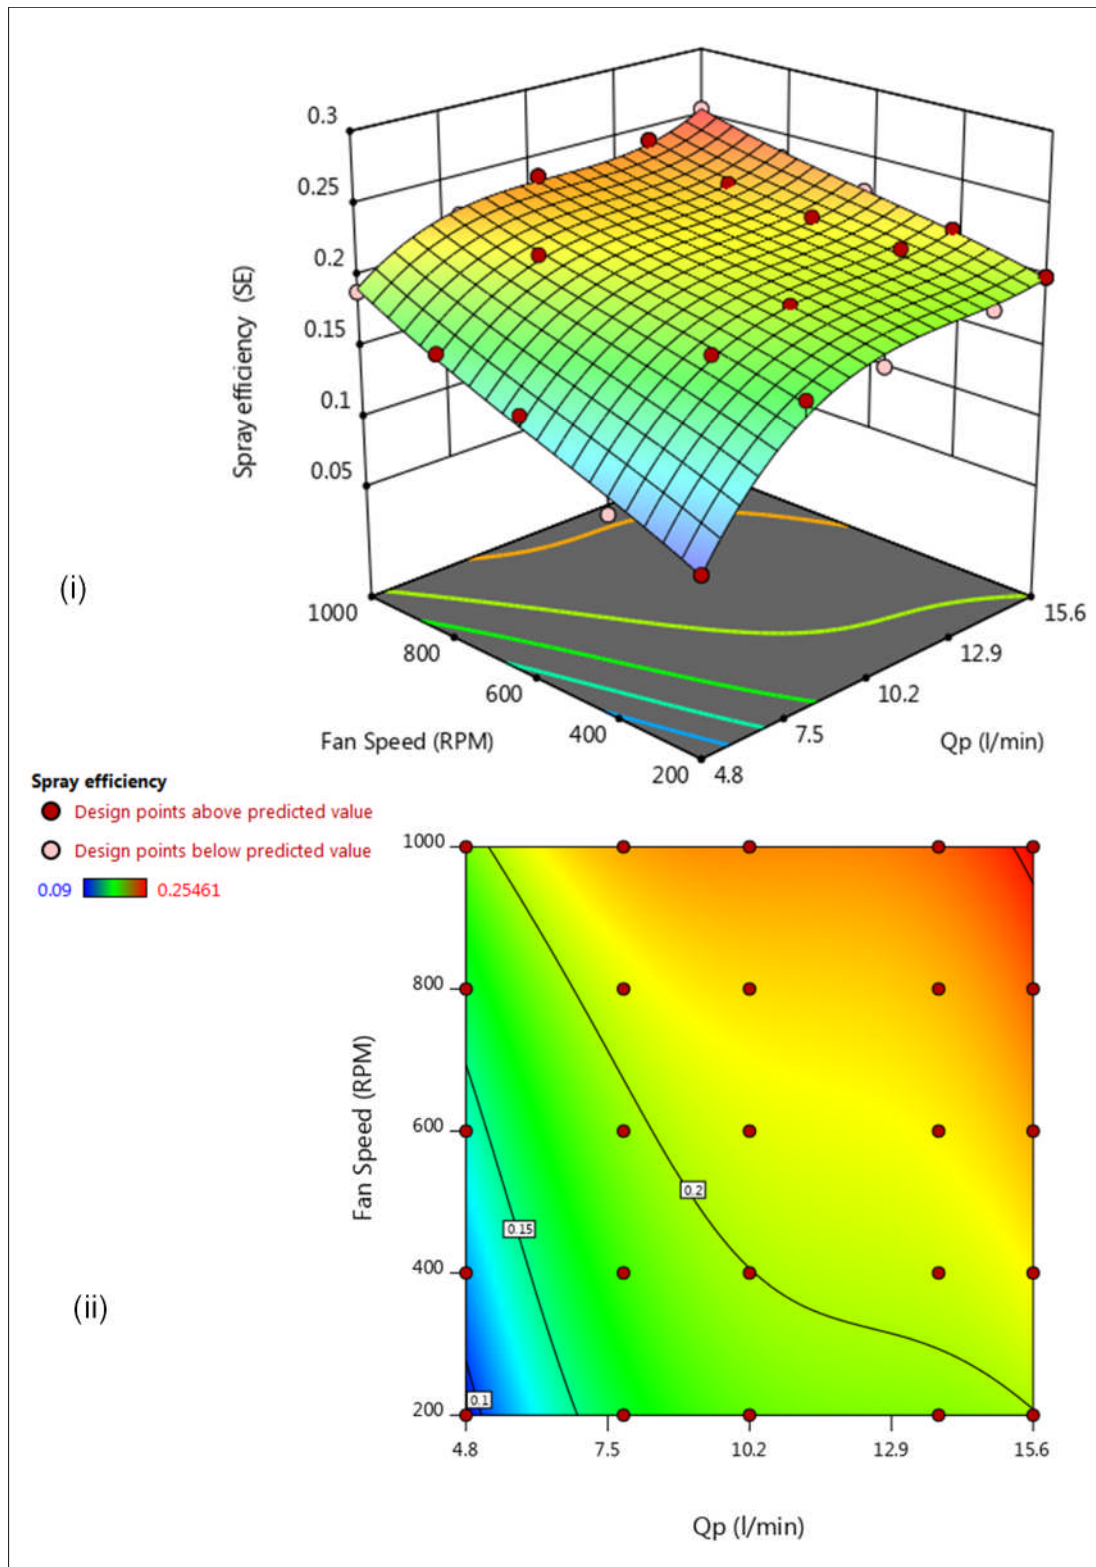
Figure 5.17 Spray efficiency as a function of the primary flow rate and fan speed, (i) shows the 3D graph of the predicted spray efficiency and (ii) shows the design points on the 2D graph.
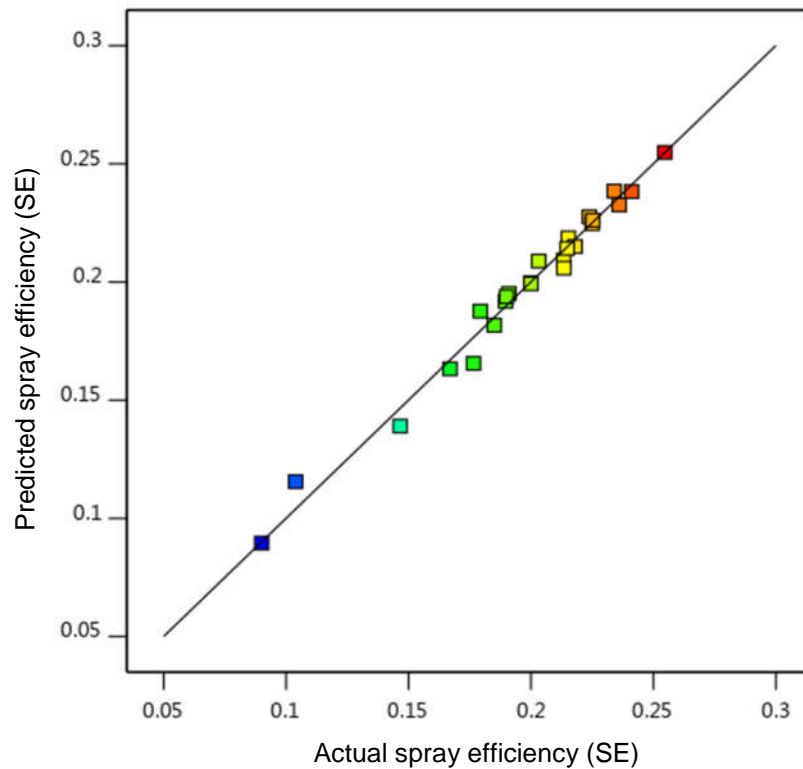
Figure 5.18 Calculated spray efficiency from equation 5.9 against the actual values calculated from the experiment.

5.4 Holistic thermal performance of the cooling system

The cooling system components are connected through each others through the cooling fluid mediums as shown in figure 5.19. Starting from the initiation of cooling at the AHU where the heat is rejected to the environment, the AHU thermal performance was found to improve with the increase of the AHU fan speed, primary loop flow rate and the spray flow rate as shown in section 5.3.2. However, the AHU is connected thermally to the remaining components of the system by the primary coolant loop where it passes through the RDHx and then to the CHx40 before been recirculated back to the AHU to reject the collected heat away from the data centre enclosure. The RDHx thermal performances was found to be improved with increasing the primary loop flow rate as shown in section 5.3.3. The CHx40 has two liquid loops where it connects between the primary loop of the AHU and the secondary loop which is in direct contact with the CPUs of the rack. The CHx40 thermal performance

was also found to improve with increasing the flow rate of both loops side as shown in section 5.3.1. Hence the increase in the primary loop flow rate is associated with the increase in the thermal performance of the three components; AHU, RDHx and the CHx40. Finally, increasing the secondary loop flow rate improves the thermal performance at the DCLC heads and as a result higher heat transfer from the chips inside the servers.
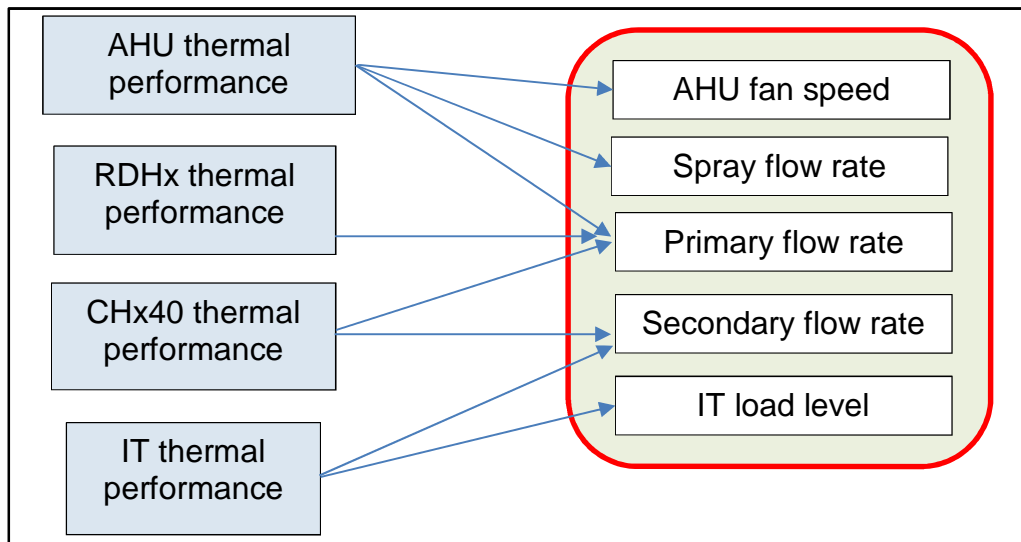
Figure 5.19 The effect of different operational conditions on the holistic thermal performance of the cooling system.

5.5 Operation conditions of the system under long diurnal runs.

Studying the impact of diurnal and weather variation on the thermal performance and power consumption of the data centre as a whole is valuable along with investigating the effect of various operation conditions of fan speeds, flow rates, and the type of cooling in the data centre facility. To study the long diurnal operation, the designed data centre facility was continuously run for three days under full utilisation of all the servers. The tests were run over different times to cover various weather conditions.

Overall, the system was run for 74 hours at winter environmental conditions of maximum and minimum ambient temperature of 11°C and 3°C, respectively. The internal and external flow rates in the system were kept constant at 15l/min and 18l/min, respectively. The AHU was kept on auto

mode to regulate the inlet temperature to the IT which was set to 22°C. To achieve this temperature, the AHU had to use different fan speeds, spray, and bypasses to regulate the required temperature as mentioned in the control algorithms in chapter 3. Temperature, flow, and power measurements were taken every minute to be analysed for the purpose of comparison as mentioned in chapter 4. The investigations were made at all levels of the system from the CPU temperatures to the AHU.

5.5.1 AHU operation

The key temperatures in the AHU at various stages of the primary loop and the ambient temperature over the three days test are shown in figure 5.20. The ambient temperature fluctuated during the three days and reached the minimum of about 3°C in the night of the third day and reached the maximum of about 10°C at the beginning of the second day. The inlet temperature to the heat exchanger (WT7) has a similar value to the temperature outlet from the IT (WT6). This is due to the fact that the heat exchanger bypass loop was kept closed. The outlet temperature from the AHU HE (WT8) is lower than the inlet temperature to the IT (WT5) which indicates that the cooling capacity of the heat exchanger is higher than the heating load as the ambient temperature is relatively low. However, the supplied temperature to the IT (WT5) is regulated to be kept almost constant at 22°C regardless of the ambient temperature. This was obtained using the bypass loop which works to keep the temperature fixed at the required level for more reliable operation of the data centre.

It can be seen from figure 5.20 that at the higher ambient temperature, the outlet temperature from the HE (WT8) approached the inlet temperature to the IT (WT5) to compensate for the increase in the ambient temperature which required less bypassing of the coolant to keep the temperature at the set value. Conversely, the large drop in the ambient temperature (point A in figure 5.20) shows a comparable reduction in the coolant temperature out of the HE which also has an effect of dropping the supplied temperature to the IT before the bypass loop manged to keep the temperature back on the track of 22°C. Finally, the fluctuation in the temperature inlet to the IT is still within the

required limit which is a result of changing the amount of the coolant that bypasses the AHU HE.
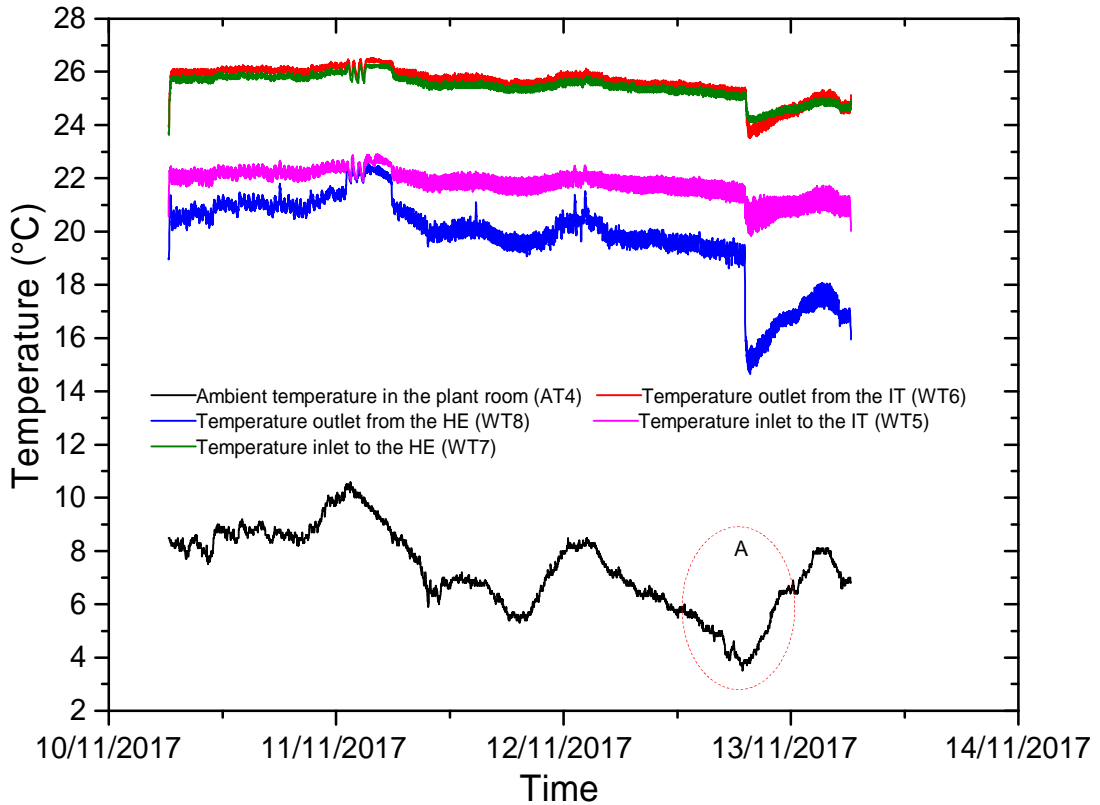
Figure 5.20 Temperature trace in the AHU for the long diurnal operation (the sensor numbers in the legend indicate the position in figure 4.1).

Figure 5.21 shows the usage of the primary loop pump, AHU fan and the spray pump for the long diurnal operation of the designed data centre facility. It can be seen that the primary loop pump speed is almost constant at about 88% of maximum pump speed to reach the target flow rate of 18l/min in the primary loop. However, there is a fluctuation of about ±2.5% which is due to the opening and closing of the bypass loop that changes the pressure drop in the loop which in turn triggers the flow rate in the loop where the pCOWeb controller must adjust the pump speed to reach the target flow rate. As mentioned previously, the fan speed is changing to achieve the required temperature assisted by using the spray when the fan reaches 80% of its maximum speed. At some points during the operation, the fan speed stayed below 80% which saves using the spray. This appears when the bypass loop is open and the water was almost fully recirculated with AHU heat exchanger

which saves the extra power required to regulate the temperature using the bypass loop and when the ambient temperature dropped which makes the high speed usage of the fan unnecessary.
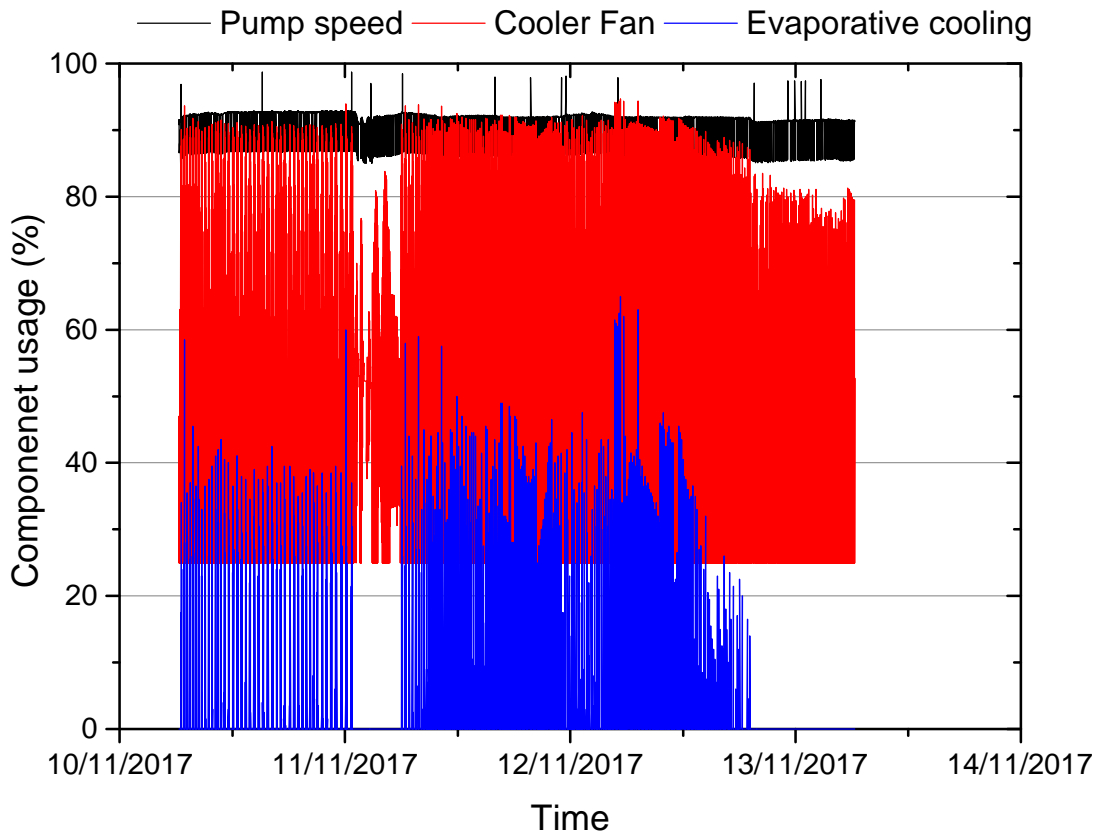
Figure 5.21 AHU fan, process coolant pump, and evaporative cooling usage during the long diurnal operation of the data centre.

The power consumption of the AHU unit is shown in figure 5.22. This includes the power consumption of all the components of the AHU collectively. It is shown that the power consumption of the unit decreases to about 35% when the fan speed dropped below the 80% of its maximum speed as the spray will not be used. This will have a positive effect on improving the pPUE and TCUE metric values of the data centre.

156
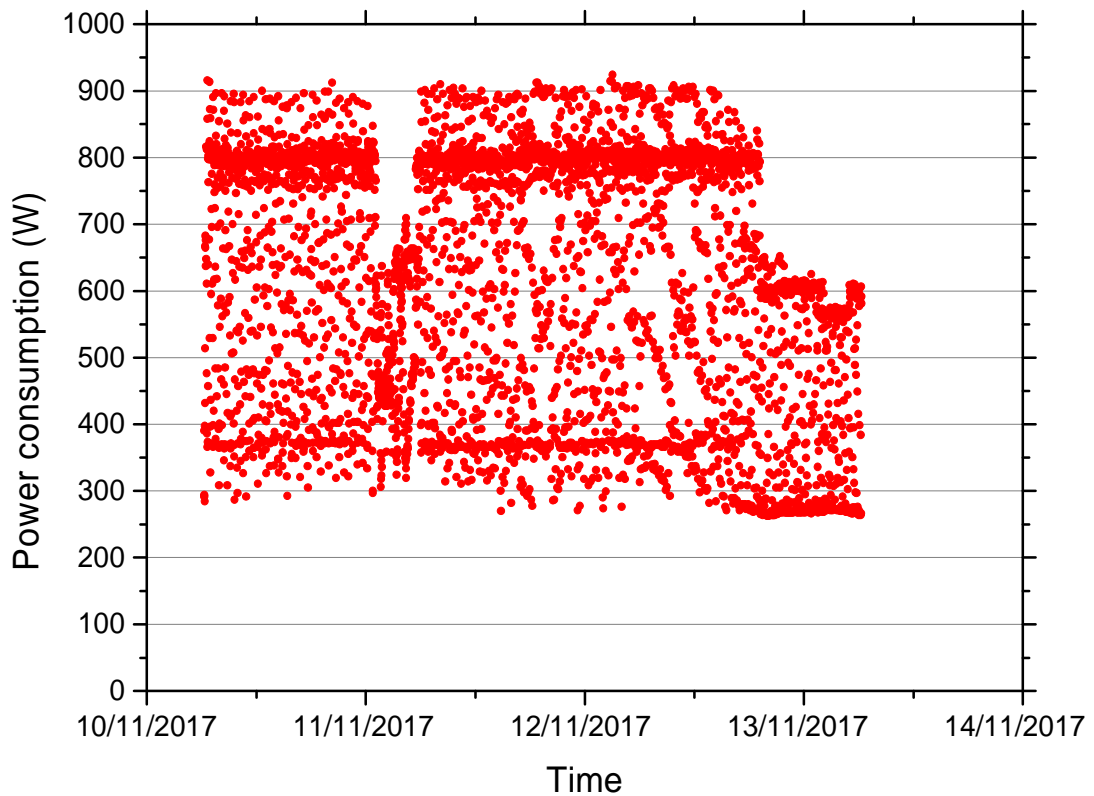
Figure 5.22 Power consumption of the AHU during the long diurnal run of the data centre.

5.5.2 CHx40 operation

The CHx40 key temperatures show a clear response to the output temperature of the AHU even if the coolant is preheated by the RDHx before reaching the CHx40 as shown in figure 5.23. The drop in the ambient temperature shows a drop of 1°C in the temperature of secondary loop coolant inlet to the rack before the AHU managed to raise the temperature back to the set value. The air temperature inside the rack represent the temperature in the top of the rack where the CHx40 reports it to calculate the dew point temperature so that the flow is controlled in case that the flow temperature reaches the dew point to prevent condensation appearing on the electrical components inside the IT.
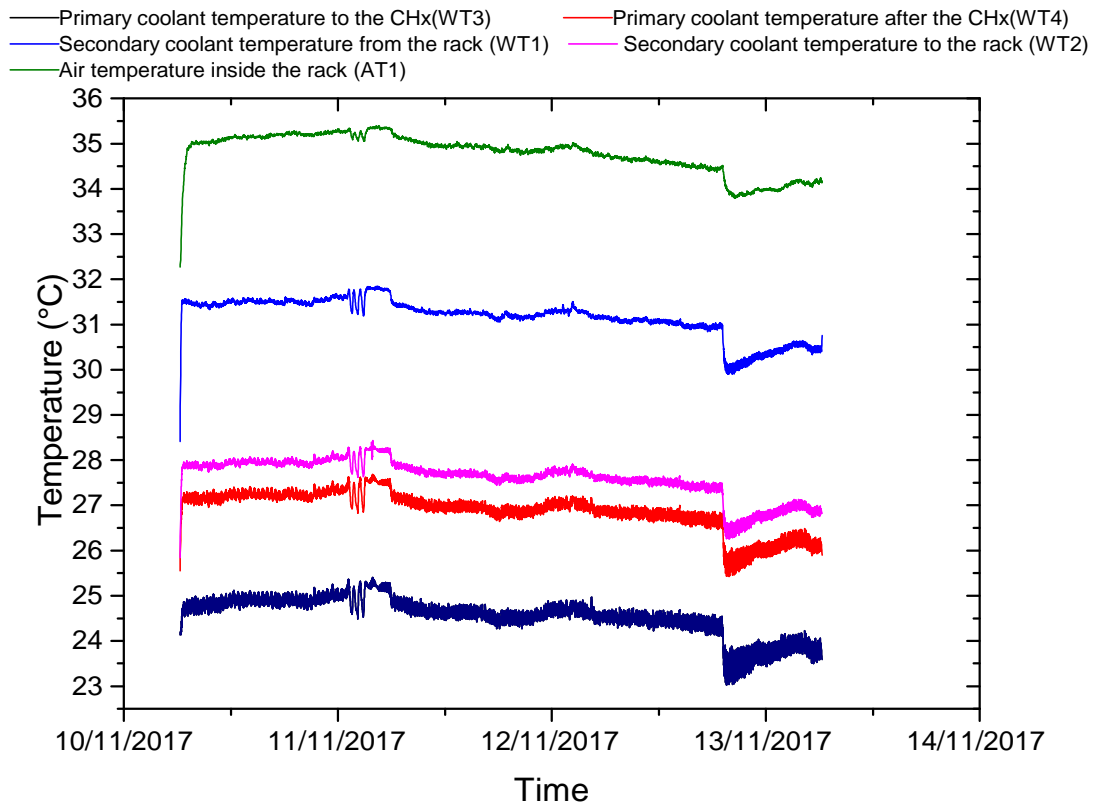
Figure 5.23 Temperature trace of the CHx40 for the long diurnal run (the sensor numbers in the legend indicate the position in figure 4.1).

5.5.3 Servers operation during the long diurnal operation

To determine the operational characteristics of the servers, different temperature points inside the servers were taken. All the thirty servers were configured to give different temperature readings of the components which are used in the analysis presented in subsequent chapters. Figure 5.24 shows the temperature measurements inside three selected servers, namely `eng06` towards the top quarter of the rack, `eng15` in the middle of the rack and `eng23` towards the bottom of the rack. The figure shows the temperatures of CPU0, CPU1, RAM0, RAM1, Ethernet, HDD, and the ambient temperature inside the servers. As mentioned previously, the CPUs are cooled by water while the other components are cooled by air. The CPUs have the highest temperature of all the components and the temperature of CPU1 within any server being always higher than CPU0 for the maximum utilisation load as they are connected in series in the cooling loop and the flow passes through CPU0

before CPU1 which means that the latter receives coolant at a higher temperature. The server fans are directly positioned to cool the RAM, HDDs, the power supply, the service processor (SP) and Ethernet. The RAM temperatures are found to have a response of 1℃ when the temperature supplied to the CPUs is dropped due to a drop in ambient conditions. This is also affected the primary loop as the air through the RAMs is cooled by the RDHx. The temperature is found to be different for the CPUs of different servers which is attributed to the flow distribution inside the rack and the variation in the pressure drop of different heads as will be investigated in the next chapter.
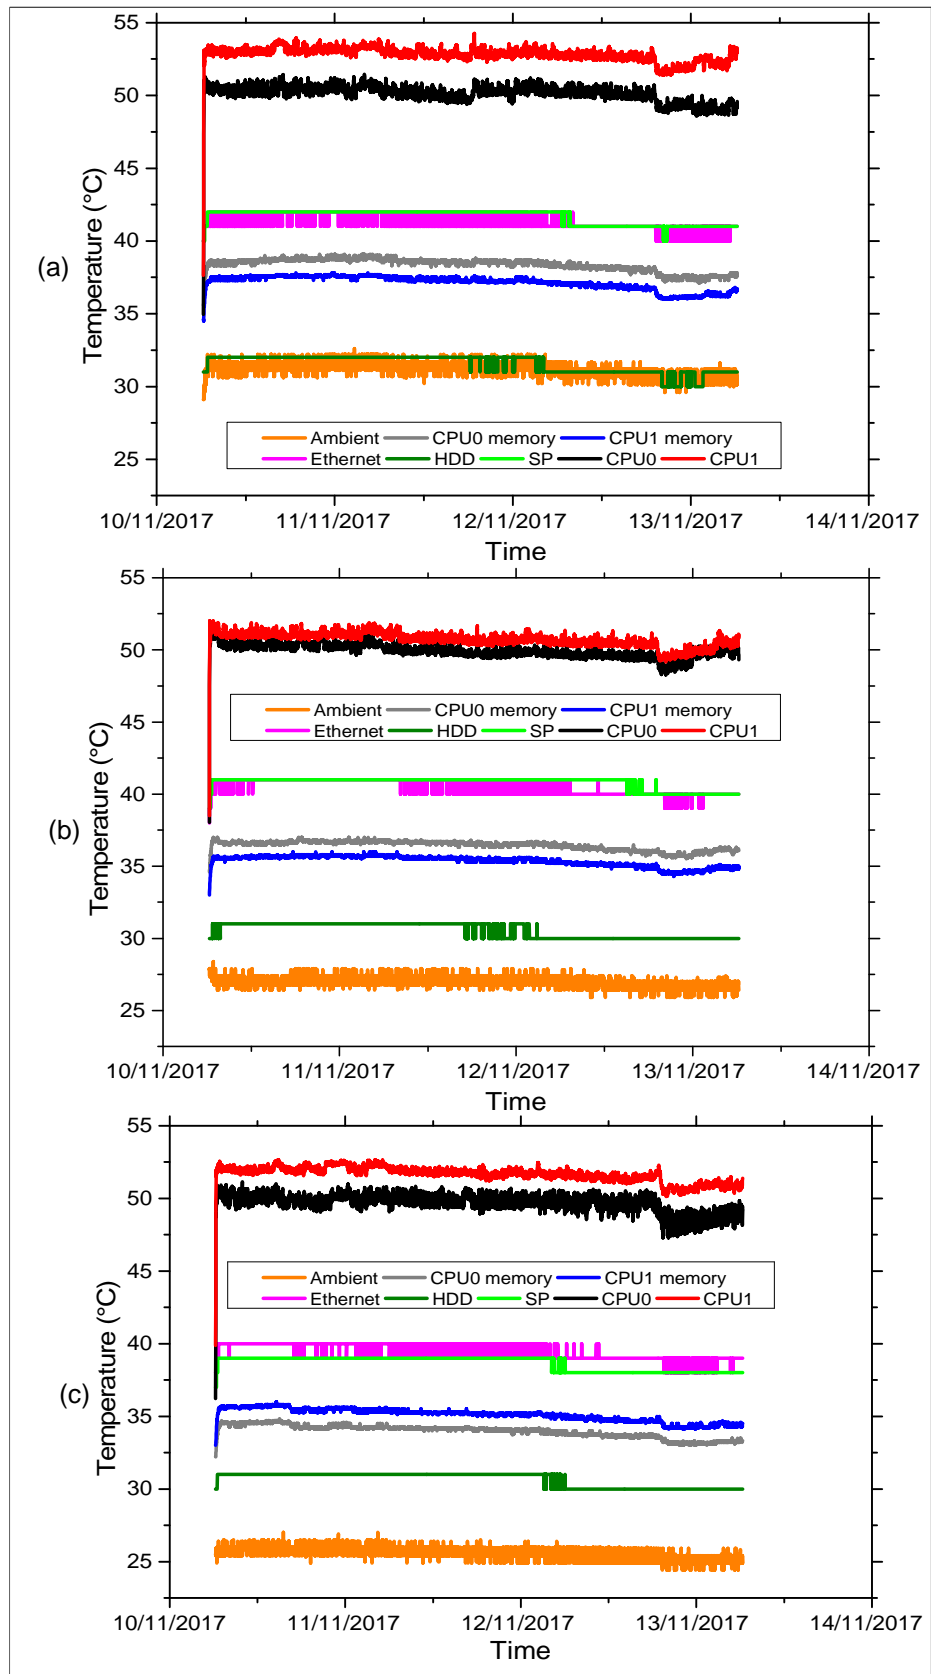
Figure 5.24 Temperature measurements inside three selected servers during the long diurnal run; (a) `eng06`**, (b)** `eng15`**, (b)** `eng23`**.**

5.5.4 Effectiveness of the data centre

The power consumption of all the components are logged and the average value of every minute is recorded to report the power consumption for the three-day experiment. The effectiveness of the data centre is reported through two metrics: pPUE and TCUE as explained in chapter 4. The instantaneous pPUE and TCUE over the long diurnal experiment for the variable weather conditions are shown in figure 5.25. The TCUE is calculated by subtracting the small fans' power consumption of the servers from the IT side based on the reported speeds of each fan in the server. As expected, both forms of effectiveness show some improvement as the power consumption of the AHU dropped when the weather temperature dropped (as shown in point A in figure 5.25) and when the AHU was able to control the inlet temperature with less use of the bypass loop (as shown in point B in figure 5.25).
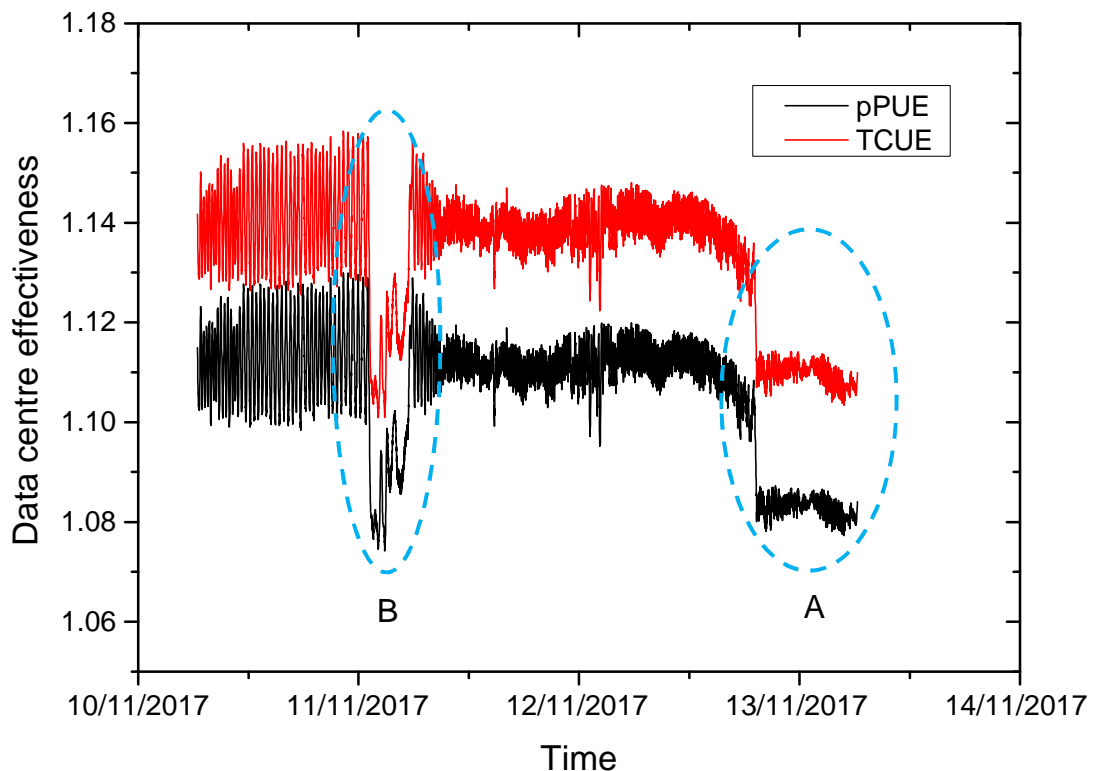
Figure 5.25 Effectiveness of the data centre as a function of time over the three days experiment.

5.6 Operational performance of the design under different IT loads

To study the operational performance of the design and examine the ability of the designed cooling system to handle the cooling load and provide a control of the temperature input to the data centre under different IT loads, a stress load scenario was proposed, as mentioned in chapter 4, to mimic a potentially real dynamic operational workload of a real data centre. In these experiments, the secondary and primary loop flow rates were kept constant while the AHU was set to auto dry mode or auto wet mode. As explained in previous chapters, the auto mode is to set the control algorithm of the AHU to control the temperatures inlet to the IT. The experiments were achieved at various ambient temperature and the comparisons with the two types of AHU operation modes were obtained under similar ambient temperatures.

5.6.1 Operational response of the AHU under variable IT load for the dry and wet auto modes.

Two multi stress runs were carried out at an average ambient temperature of 13℃ over two consecutive days. The AHU in the first test was set to auto dry mode while in the second day it was set to auto wet mode. The tests were initiated with a secondary flow rate of 6l/min and primary flow rate of 18l/min. The set point temperature was maintained at 20°C during the 7 hours of experiments so that the AHU had to regulate the temperature of the data centre with both modes of dry (figure 5.26) and wet (figure 5.27) conditions. It should also be mentioned that these two experiments were carried out before installing the RDHx so that the AHU is only cooling the CPUs of the servers using the DCLC heads.

The rack workload profile has essentially four types which includes different stress levels over a period of 7 hours:

1. Fluctuating load over 120 minutes. In this load, the first hour includes pulsation of 5 minutes full utilisation and 5 minutes idle while the second hour includes 15 minutes full utilisation pulsation followed by 15 minutes idle operation.

2. Idle operation of the rack for one hour where the servers work on 0% utilisation before the rack is stressed again.
3. Moderate load over 120 minutes. The first hour of the moderate load is to stress the full rack at 25% and the second hour is to stress the rack at 50%.
4. High load for 120 minutes. The last two hours are with the rack stressed to 75% load for an hour and then at the full load of 100% utilisation for the second hour.

Figures 5.26 and 5.27 show the operational history of the AHU and the power consumption over the 7 hours of tests for the dry and wet operation, respectively. The ambient temperature is in the same range for the two days. The AHU managed to regulate and keep the inlet temperature to the IT (to the CHx40 in these two experiments) constant for the different types of loads. However, as expected, the power consumption of the AHU varies depending on the load as well as the weather conditions and as a result the effectiveness metrics of the data centre vary.

In the first hour of the test, where 5 minutes fluctuation tests were carried out, the stress pulses result in fan speed pulses which cause an increase in the power consumption of the AHU. In the wet experiment, the spray is activated when the fan speed reaches 800 rpm, which results in an additional power consumption as shown in figure 5.27. However, this is only shown for very short times of less than a minute as the load is not continuous. For the second hour of the test which consists of 15 minutes of load pulses, the AHU unit responds to the load in a similar way but the fan keeps ramping to high speeds as the load is imposed for longer. However, in the wet experiment, the power consumption is again higher than the dry experiment (figure 5.26) due to the activation of the spray.

In idle operation of the rack at the third hour of the test, the fluctuation in the fan speed is relatively small compared with the other stress scenarios to regulate the temperature. It can also be seen from figures 5.26 and 5.27, that the power consumption of the unit is relatively similar for both dry and wet operation of the AHU unit.

In the first hour of the moderate load (the fourth hour in the test), the rack is stressed for 25% utilisation. The power consumption of the unit is approximately 15% higher than under idle operation while the fan speed stayed beyond the point where the spray is needed which is shown from the two figures 5.26 and 5.27. On the other hand, in the 50% utilisation stress in the fifth hour of the test, the power consumed by the AHU was 30% higher than the idle hour test for both dry and wet experiments. However, it is shown in figure 5.27 for the 50% utilisation test that the power consumption of the unit ramped up at three times very briefly while at some points it reached even less than the dry case as of using the spray which resulted in using less fan speeds for short times after activating the spray.

Finally, for the high IT load of 75% and 100% utilisation, the average power consumption of the AHU is almost the same for both stress levels in the dry and wet cases. However, the power consumption in the wet case is higher than in the dry case by about 16%. This is due to the fact that the heat generated by the servers is low since only the CPUs heat is captured by the system while the remaining heat generated by the servers is transferred to the room and becomes part of the cooling load of the building and not part of the cooling load of the AHU. However, the AHU managed to deliver almost stable temperature at 20° for both cases of dry and wet operations. In the wet operation and as the fan reaches the 80% threshold, the spray is activated even if the cooling load is low which increases the power consumption peaks.
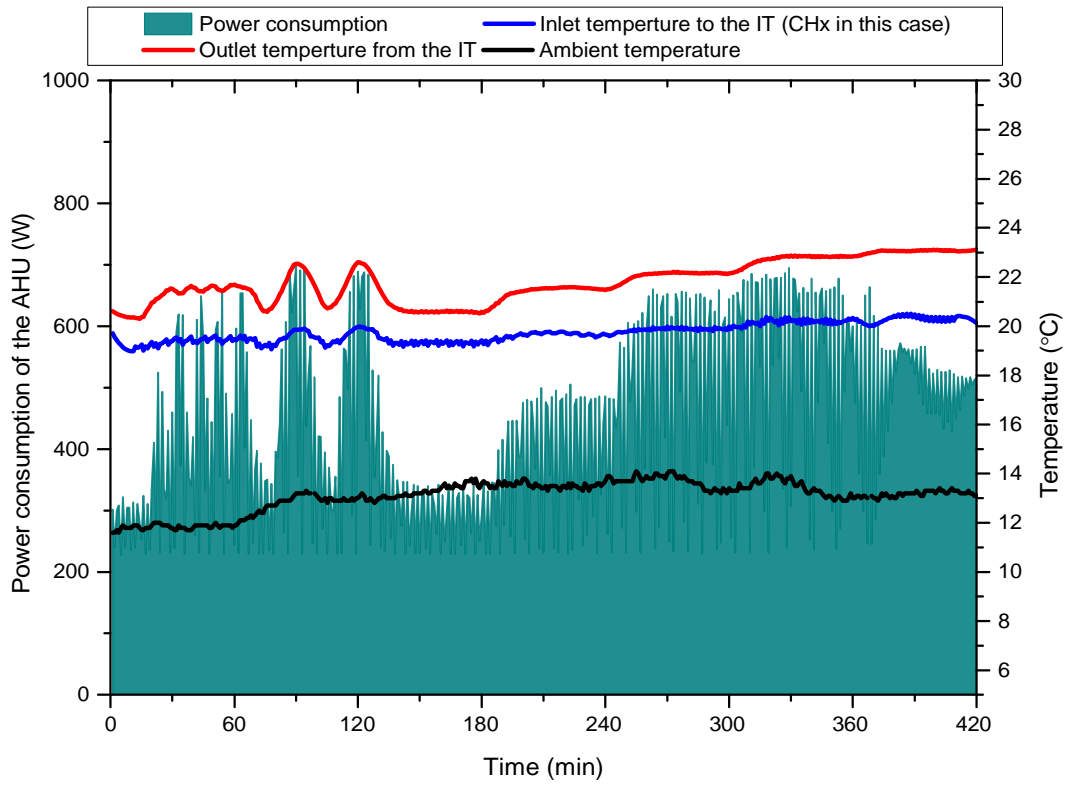
Figure 5.26 Operational characteristics of the AHU in the dry case.
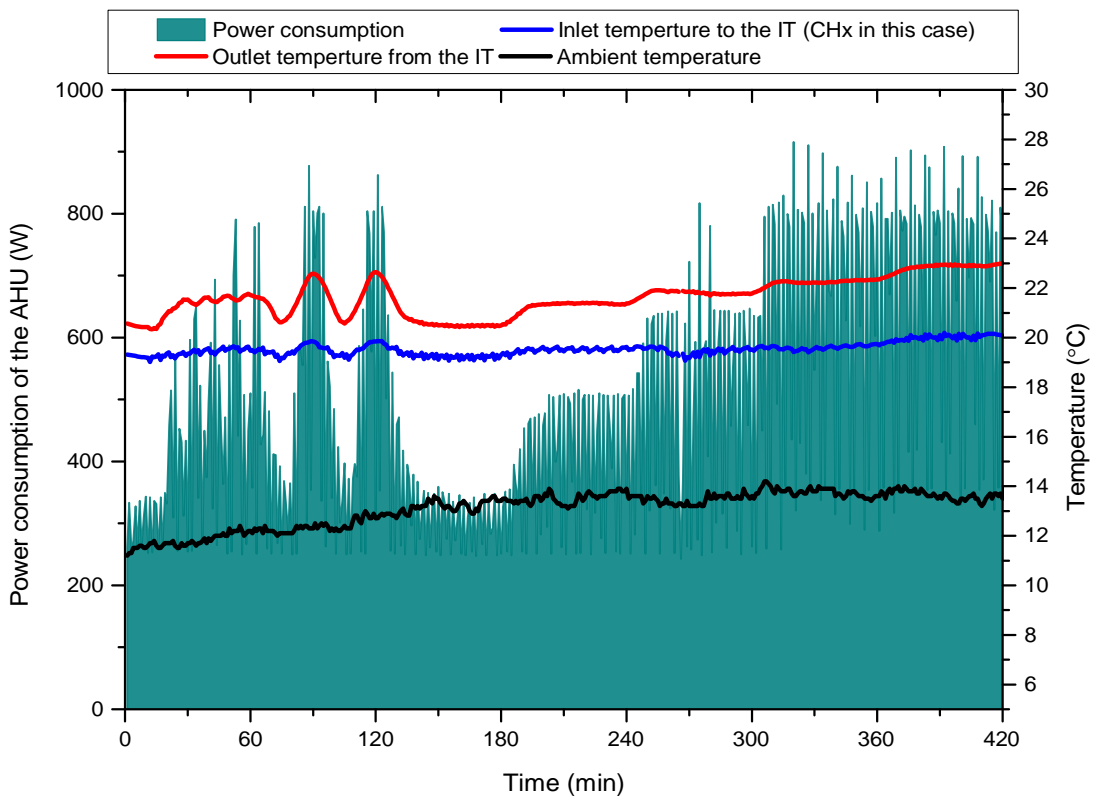
Figure 5.27 Operational characteristics of the AHU in the wet case.

5.6.2 Operational response of the designed cooling system under variable IT load for different ambient conditions.

5.6.2.1 AHU operation

Further sets of experiments were carried out with the same stress scenarios as in the previous section. However, these experiments reveal the effect of ambient temperature on the stress level during these tests. It was difficult to obtain stable ambient conditions for seven hours in any day. For this, the first experiment was achieved at relatively higher ambient conditions (between 12.5°C and 14°C) and the second experiment was achieved at lower ambient conditions (between 9.5°C to 12°C). It should be mentioned that these experiments were carried out after installing the RDHx which means that all the heat generated by the servers is captured by the cooling system and transferred to the environment by the AHU.

Figures 5.28 and 5.29 show the temperature characteristics in the AHU for the relatively higher and lower ambient conditions, respectively. It is shown in both figures that there is a lag of about 10 minutes between the start of the fluctuation stresses and the response of the AHU. It is also shown that the outlet temperature from the IT (WT6), the inlet temperature to the HE (WT7) and the outlet temperature from the HE (WT8) are all affected by the changes in the IT stress level and the ambient conditions. However, the supplied temperature inlet to the IT (WT5) is more stable to the changes of stress level and ambient conditions as a result of the bypass loop temperature regulation control. From the comparison of the two figures, it is clearly seen that the AHU control algorithm managed to control the inlet temperature to the IT for the different IT loads and ambient conditions.
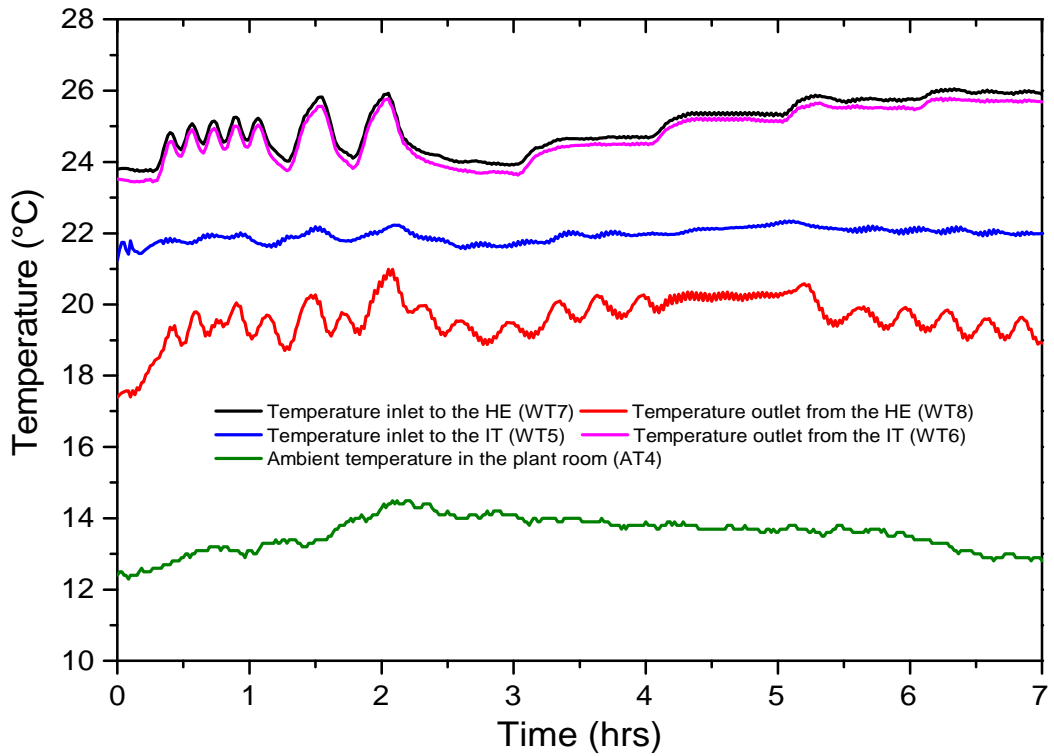
Figure 5.28 Temperature characteristics of the AHU at relatively higher ambient temperature (12.5°C and 14°C).
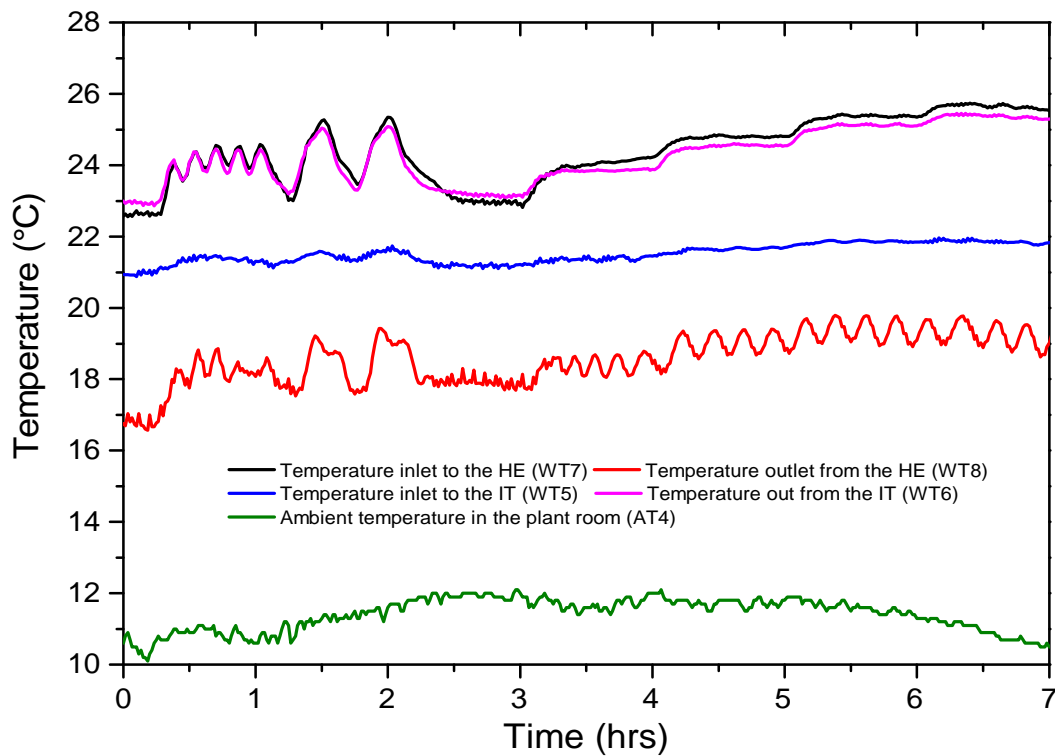
Figure 5.29 Temperature characteristics of the AHU at relatively lower ambient temperature (9.5°C to 12°C).

Figures 5.30 and 5.31 shows the AHU fan speed, process water pump speed, spray pump speed, and the utilisation for the seven hours experiments of the relatively high and low ambient temperatures, respectively. The process water pump is working in a limited range between 85% and 93% of its maximum speed to achieve the required flow rate of 18l/min. As discussed previously, the small variation of the process water pump speed is due to the process of opening and closing the bypass loops to regulate the temperature where the AHU HE is connected and disconnected from the loop which is altering the pressure drop in the system and hence the pump must adjust its speed to achieve the targeted flow rate. The fan speed is following the control algorithm to regulate the temperature and it can be seen in both figures that the fan speed is linked to the stress level where it is fluctuating in the first two hours for the fluctuation loads, reduces in speed for the third hour of the idle operation followed by a step increase in the average speed of each hour following the load increase until the end of the experiment. The spray pump was activated twice in the fluctuation load at the higher ambient temperature experiment where the fan speed exceeded 800rpm and twice in the high load hours (the 6[th] and 7[th] hours of the experiment) where the fan ramped up to provide the required inlet temperature to the IT.
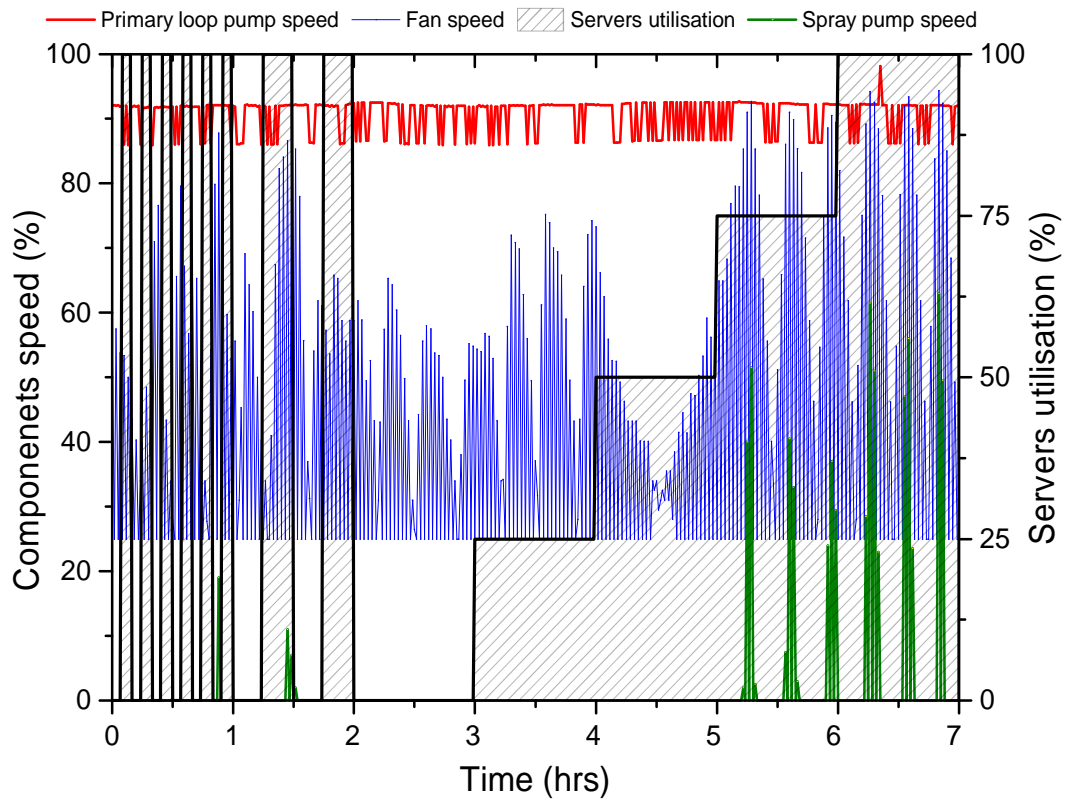
Figure 5.30 AHU components usage during the 7 hours operation at the relatively higher ambient temperature (12.5°C and 14°C).
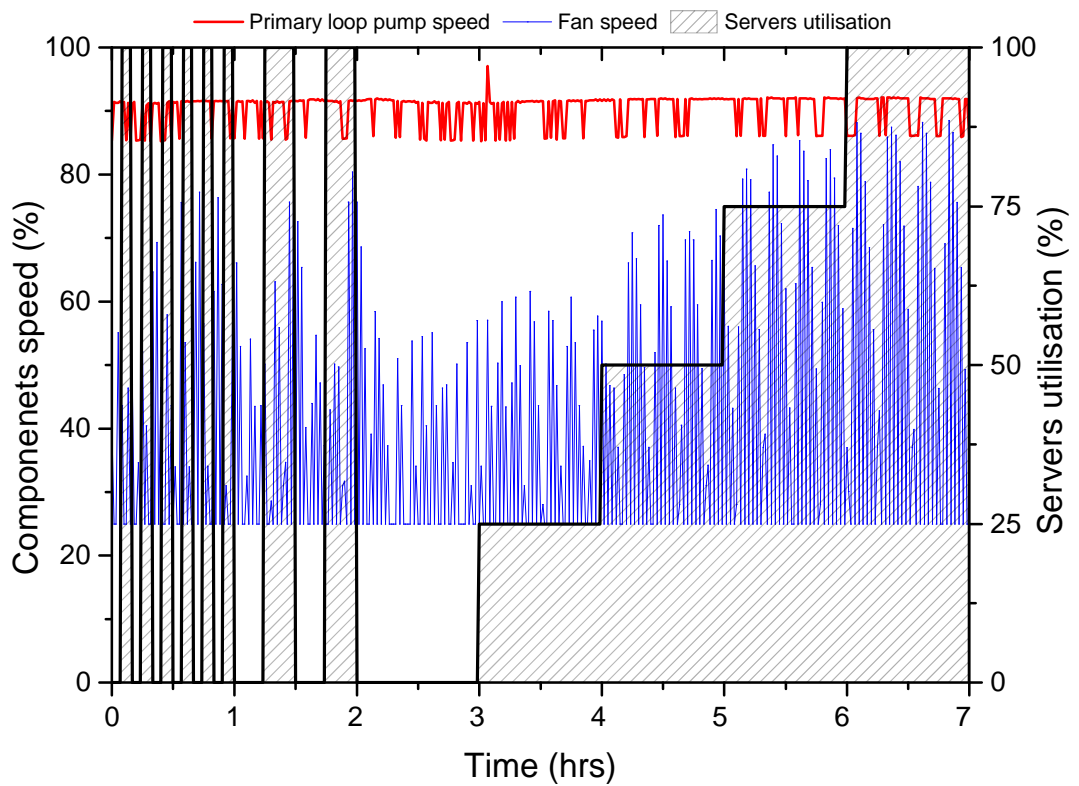
Figure 5.31 AHU components usage during the 7 hours operation at the relatively lower ambient temperature (9.5°C to 12°C).

The total power consumption of the AHU in the two experiments is shown in figures 5.32 and 5.33, respectively. The effect of the ambient temperature is shown clearly on the power consumption of the two experiments. In general, the power consumption of the AHU at the cold ambient conditions is less than the power consumption at the higher ambient temperature for the seven different types of rack utilisation loads. At the low ambient temperature experiment (figure 5.33), the power consumption at the 5 minutes fluctuation experiment never exceeded 600 W while for the high ambient temperature (figure 5.32) this happened occasionally. The difference in power consumption is also clearer at the 25% utilisation load where the power required to keep the temperature at the targeted level is about 500W in the high ambient temperature experiment while it is about 380W for the low ambient temperature experiment.
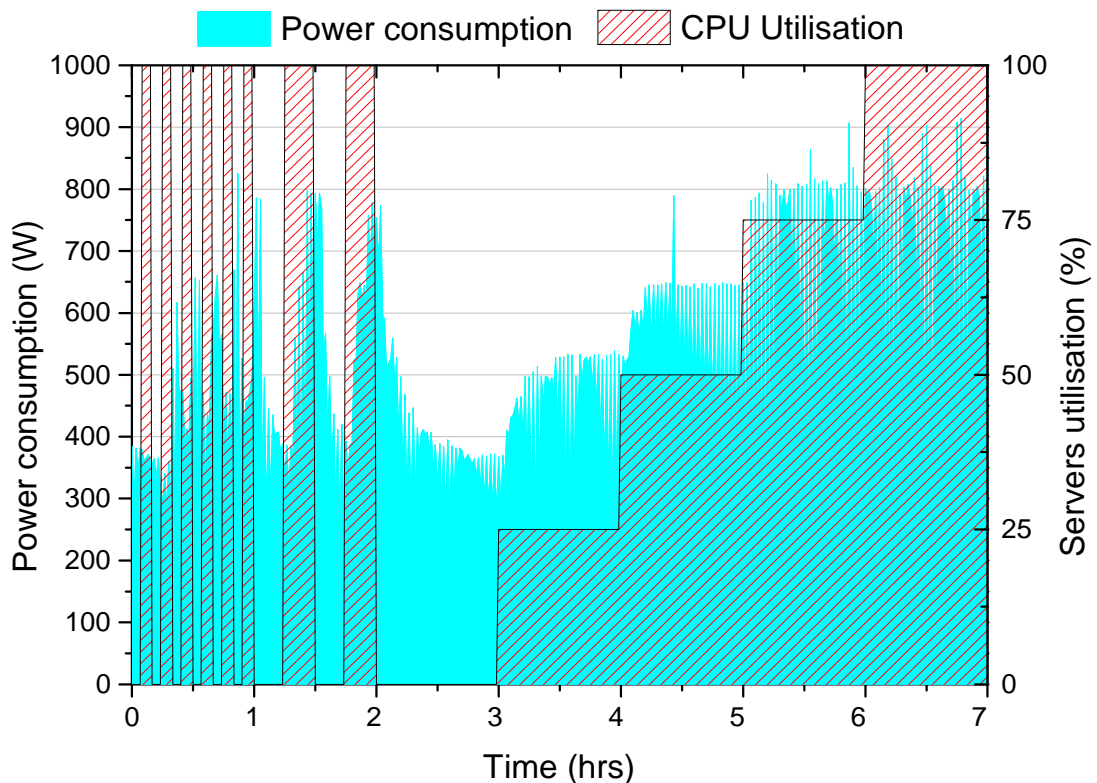
Figure 5.32 Power consumption of the AHU at the relatively higher ambient conditions (12.5°C and 14°C).
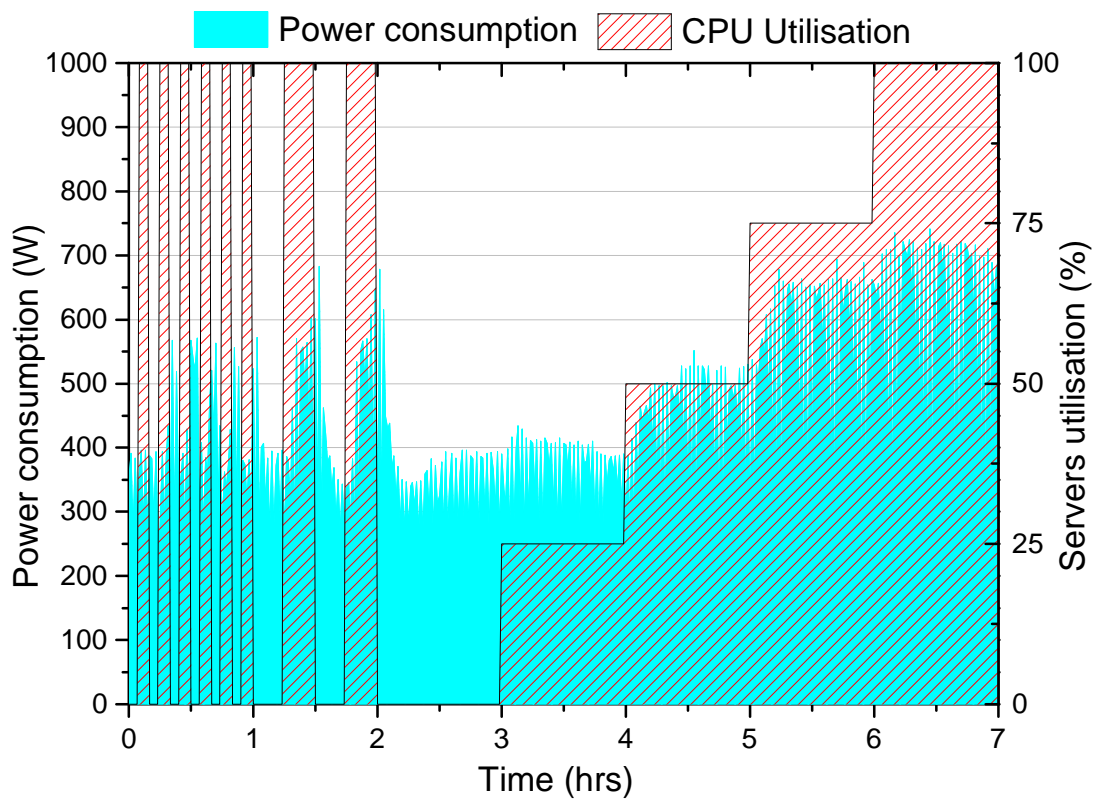
Figure 5.33 Power consumption of the AHU at the relatively lower ambient conditions (9.5°C to 12°C).

5.6.2.2 CHx40 operation

The response of the CHx40 to the 7 hours multi-level stress test is shown in figure 5.34. As the inlet temperature to the IT (WT5) is the same in the high and low ambient temperature experiments, only one of the experimental results of the CHx40 will be presented and discussed here. The power consumption of the CHx40 is almost the same for all the experiment as the pumps speed are kept constant. As mentioned in the previous section, the AHU keeps the inlet temperature to the IT constant (within $\pm 1°C$) where the process water passes through the RDHx, which gains some heat, before reaching the CHx40 in the primary loop. It follows that the inlet temperature to the CHx40 in the primary loop (WT3) is also constant and follows the same responses of WT5 in the AHU but with a higher temperature which represents the heat rejected by the air-cooled components of the servers.

The coolant outlet temperature from the CHx40 in the primary loop (WT4), and the inlet (WT1) and outlet (WT2) of the coolant temperatures in the secondary

171

loop all follow those of the 7 stress level scenarios. The ambient temperature inside the rack is measured by the CHx40 sensor, which is also used to determine the dew point temperature to be used in the control system to prevent condensation inside the IT equipment. It is also measured during the 7 hours experiment and shows a direct relationship with the rack load.
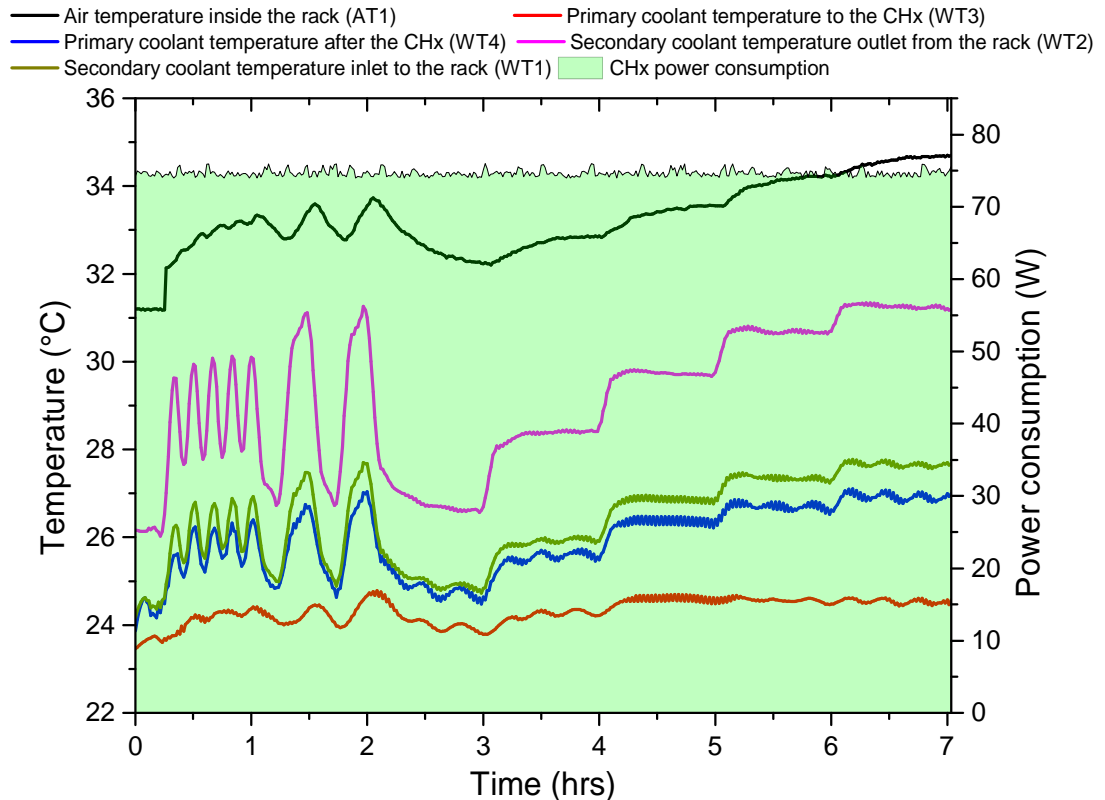
Figure 5.34 Temperature and power trace of the CHx40 for the multi-stress level run.

5.6.2.3 Servers operation

The temperature measurements of the different components of the servers are shown in figure 5.35 for two selected servers (`eng06` and `eng23`) for the seven hours multi- stress level experiments. The figure shows the responses in temperature for the different stresses. It can be seen that the CPU temperatures raise with the applied stress. However, the other components maintain a steady temperature during the experiment.

172

As mentioned previously, every server has two CPUs (CPU0 and CPU1) which are connected in series in the cooling loop which means that CPU0 receives the coolant with lower temperature than CPU1. It is shown in figure 5.35 that the temperature of CPU0 is lower than CPU1 for the 1st, 2nd,3rd and 7th hours of the experiments. However, for the remaining stress levels (4th, 5th, and 6th) the CPU1 can have lower temperature than CPU0. The main reason for this is that at 100% utilisation of the servers, both CPUs of each servers are stressed to maximum percentage so that they generate the same amount of heat. Stresses lower than 100%, the controller of the CPUs distribute the load arbitrarily, and CPU1 can carry less load than CPU0 and vice versa which affects the temperature distribution of the CPUs. This is shown in figure 5.35 for the two selected servers at the 4th and 6th hours of the experiment.

To include the CPUs temperature in the analysis and investigate the effect of different operation characteristics on the servers operation, the average CPU temperature (CPU0 and CPU1) is considered across all the servers. Figure 5.36 shows the CPU temperatures for all the rack under the seven hours experiment for the multi-stress levels. The CPU temperature of the servers are different from each other depending on many factors such as CPUs head pump assembly, thermal paste, and server position. This will be investigated in detail in chapter 6.
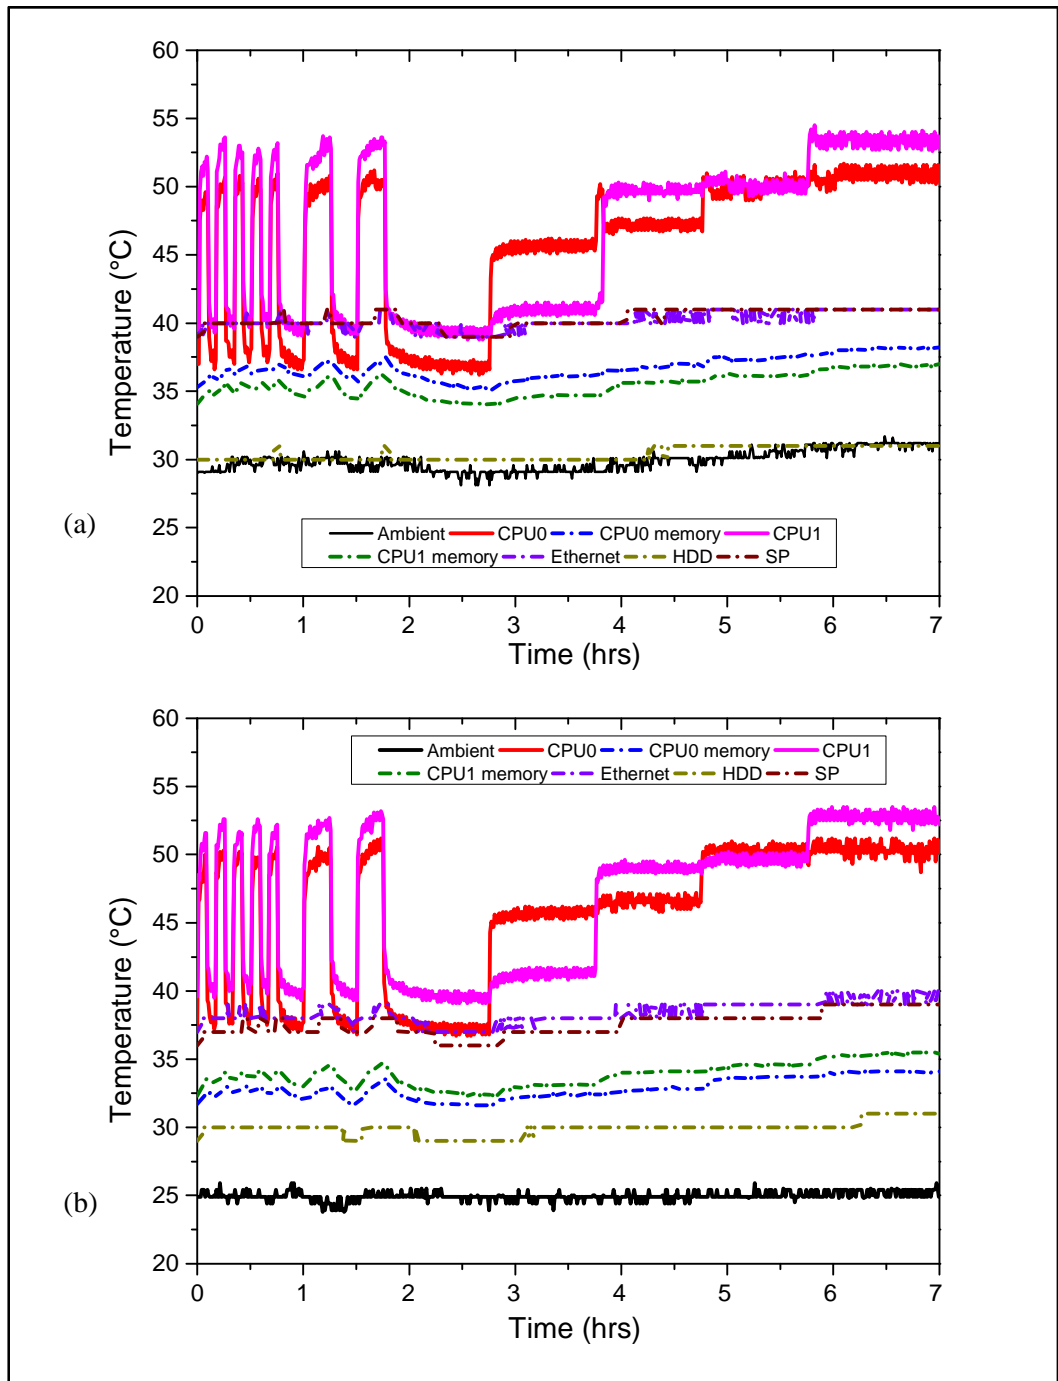
Figure 5.35 Temperature measurements inside three selected servers for the multi-stress level run; (a) `eng06` **and (b)** `eng23`**.**
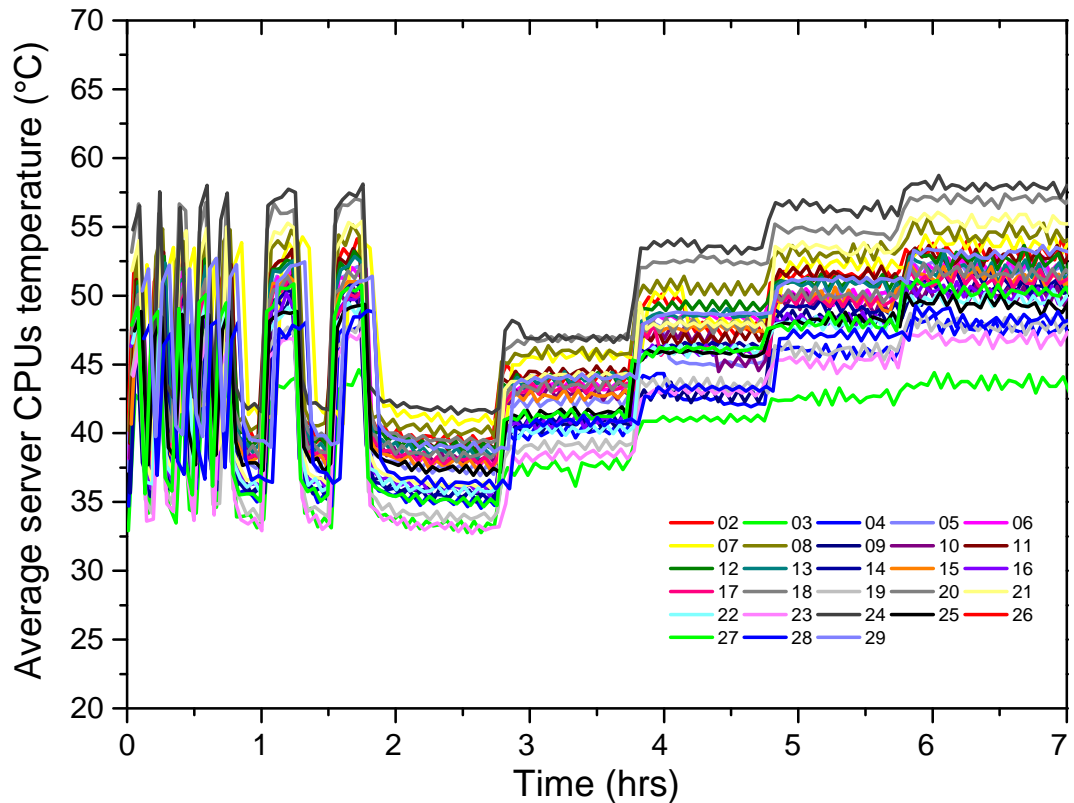
Figure 5.36 Average CPUs temperature for all the servers of the rack under the multi-stresses level experiment (the numbers in the legend represent the servers number in the rack).

5.7 Summary

In this chapter, the designed cooling system was examined and characterised to quantify the potentials for energy savings and provide a reliable means of cooling the DCLC chiller-less data centre. Therefore, the dynamic response of every component was reported and analysed in terms of thermal and power consumption characteristics. The long operation of the system for various ambient conditions and IT loads was also reported in terms of the key temperatures of the components, power consumption and effectiveness of the whole system from the chip level to the facility level.

The pressure drop in the AHU and across its heat exchanger is important as higher pressure drop means higher power consumption and less thermal benefits. The pressure drop across the heat exchanger of the AHU was found to be increased by 33% when increasing the fan speed from 250 to 1000 rpm. Utilising spray evaporation in the AHU should be considered when selecting

175

the proper fan for the design as it was found that the pressure drop increases by 145% when the spray flow rate is 30l/min.

The AHU fan power consumption constitutes the largest portion of the power consumption of the cooling system and increases dramatically in the high speeds of more than 600rpm where it consumes 960W at the speed of 1000rpm. The second largest contribution to power consumption is the spray pump where it consumes up to 300W at the maximum speed. The high power consumption of the spray pump is due to the high pressure drop in the nozzles and the solenoid valves. This is follows by the power consumption of the process water pump of 250W at the highest speed.

Activating the spray should be weighed up against the thermal benefits which are obtained from the latent heat of evaporation. Using the spray found to reduce the thermal resistance of the AHU heat exchanger by 50%. However, the amount of water that evaporates is only 25% of the total sprayed at the maximum air flow rate and primary water flow rate. Therefore, a proper design for the AHU is necessarily to collect the unevaporated water to be reused. It is also important to not over spray the heat exchanger with unnecessary amounts of water, which adds direct costs of pump power consumption and indirect costs associated with water usage.

The designed cooling system managed to provide stable coolant temperature to the IT regardless of the ambient conditions and the IT load. However, there is a lag of about 10 min between the actions produced by the servers and the response of the AHU. Therefore, it is important to consider this time lag in future designs of data centre cooling to be reduced.

Finally, it was observed that CPU0 of the servers in the rack generally had lower temperature than CPU1. This is justified by the fact that CPU1 receives higher coolant temperature than CPU0. However, CPU1 can have lower temperatures than CPU0 in some cases depending on the IT load distribution between the CPUs of a partially loaded server. Moreover, it was spotted that the CPUs of different servers in the rack have different temperatures for the same IT load.

Chapter 6 Flow distribution and pumping configuration inside the DCLC rack

6.1 Introduction

In this chapter, the thirty liquid cooled servers at the rack level are tested with regard to the coolant flow rates, pumping configurations, and flow distribution. The effect of these parameters on the CPU temperatures, power consumption, effectiveness, and efficiency of the rack is investigated. The study is also extended to optimising the rack design through examining flow maldistribution. The results of this chapter address the gap in the research of characterisation of the flow distribution in liquid cooled servers at the rack level as well as exploring the effect of the flow rate and pumping configuration on the IT.

As explained previously in chapters 3 and 4, the pumping configuration in the rack is either central or distributed. With central pumping, the two large pumps (connected in series) used to pump the coolant inside the rack and through the CPU cooling units, with distributed pumping, the pumps inside the servers at the CPU heads are running in addition to the central pumps. It should be mentioned that the pumping is purely a distributed model only when the pumps in the CHx40 are deactivated. However, in this study, the term "distributed pumping" refers to the configuration state when the servers' pumps are activated while the CHx40 pumps are either switched on for different speeds or off.

The objectives of this chapter are:

1. The flow distribution and the effect of this on the CPU temperatures of the servers will be investigated.

2. Optimisation of the design to obtain uniform flow distribution inside the servers.

3. The effect of the flow rate and flow configuration on the CPU temperatures under different IT loads will be studied.

4. The effect of these parameters on the effectiveness of the data centre including the pPUE and TCUE under different pumping configurations will be investigated.

6.2 Flow distribution

As discussed in section 4.10 of chapter 4, the uniform flow distribution in the rack for all the servers is necessary for reliable operation of the IT equipment. Experiments show that there is a high variation of the CPU temperatures of the different servers in the rack for the same IT load as shown in figure 5.35 in chapter 5. It was also noticed that some of the servers reach the critical CPU temperature while other servers still work in the safe CPU temperature range. The CPU temperature variation was attributed to the flow maldistribution due to the pressure drop in the manifolds, differences in pressure drop across servers' modules, the thermal interface material between the servers' module and the CPU surfaces, and the way of assembling the server module over the CPU. The present study focuses on the flow maldistribution due to the losses in the manifolds.

To study the flow distribution, EPANET open source software is used to analyse the flow in the rack for the central and distributed pumping under different CHx40 pumps speeds. To simulate the cooling loop system, every component in the loop is drawn and represented as shown in figures 4.5 and 4.6 in chapter 4. The software is capable of representing the pipe segments, tanks, distribution nodes, and pumps. The properties of each component are fed into the software such as the length and diameter of the pipes, elevations, minor and major losses, tank properties and different pump curves for different speeds. The full methodology and the boundary conditions of the simulation have been discussed in chapter 4.

6.2.1 Model validation

The models of the rack level flow network are validated against the experimental results for the two configurations of the central and distributed pumping. This is also achieved across five CHx40 pump speeds and also for the case where the CHx40 pumps are deactivated.

Figures 6.1 and 6.2 shows that the EPANET software model of the flow inside the rack agree with good accuracy when compared with the experimental results. The average percentage error between the experimental and calculated results for the central and distributed pumping is 2.3% while the maximum error is 3.5% for the central pumping and 5% for the distributed pumping.

The figures represent the relationship between the total coolant flow rate and the pump speed. As mentioned in chapter 4 (refer to figure 4.7), the installed central (CHx40) pumps have five different speeds with pump speed 0 representing the point where both of the pumps are off. For the central pumping case, there is no flow in the loop when the CHx40 pumps are off (Pump speed 0) as the servers' modules pumps are deactivated. However, in the distributed pumping configuration, there is a flow rate of about 7 l/min when the central pumps are off.

In both cases, figure 6.1 and 6.2, the benefit of increasing coolant flow rate between pump speed 4 and 5 are small compared with the other pumps speeds. This can be attributed to the pumps operation curve as shown in figure 4.7 (refer to chapter 4) where the pump curves are very close to each other between pump speed 4 and 5 beyond the flow rate of 7.5 l/min.
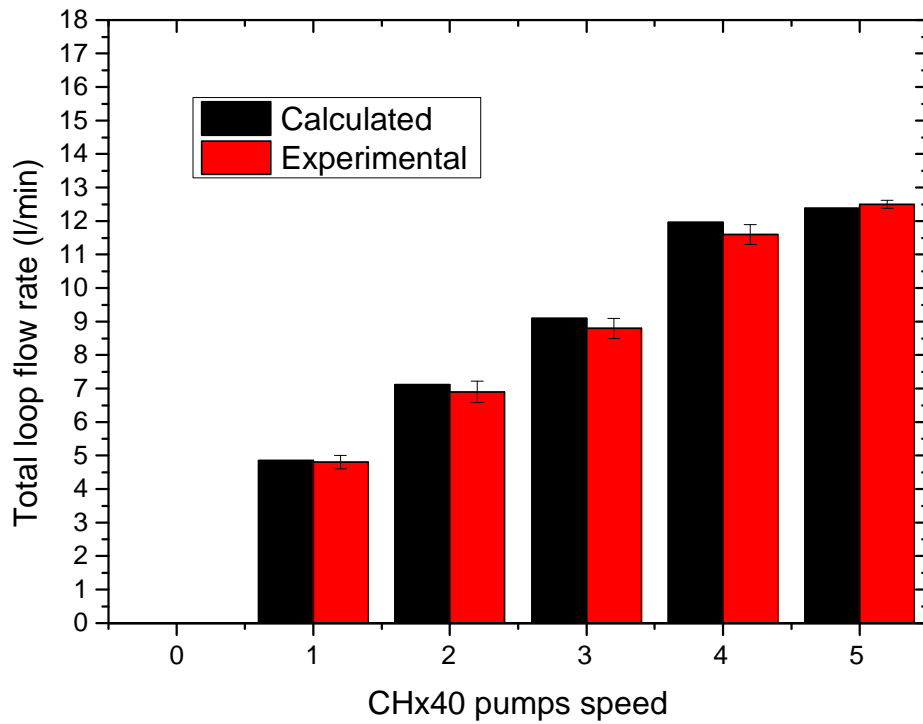
Figure 6.1 Validation of the calculated loop flow rate using the EPANET software against the experimental data for the central pumping case.
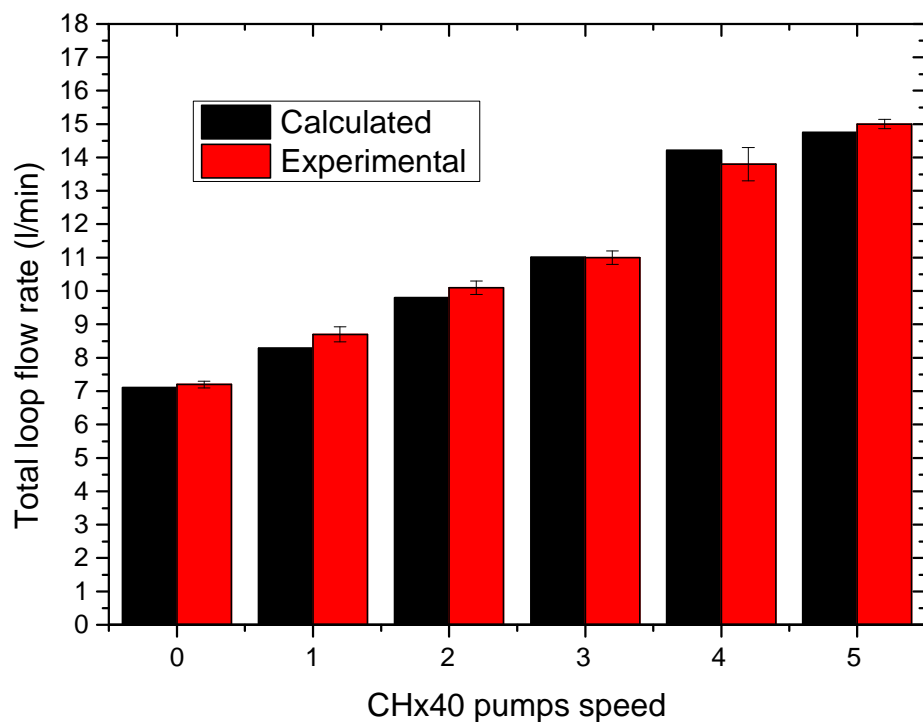
Figure 6.2 Validation of the calculated loop flow rate using the EPANET software against the experimental data for the distributed pumping case.

6.2.2 Centralised pumping

The calculated flow distribution in the rack of the thirty servers is presented in this section for the centralised pumping configuration. The flow distribution is shown in figure 6.3. It can be seen that the flow through servers at the top of the rack are larger than the ones at the bottom of the rack. As mentioned in chapter 4, the calculations assume that all the servers' modules have the same pressure drop resistance which is due to the major and minor losses in the servers' modules caused by the stäubli valves, the flexible corrugated hose and the microchannels of the heat sink. As a result, the resulted flow maldistribution is only caused by the frictional losses in the supply and collection manifolds of the rack. These losses are caused by the friction between the coolant and the internal manifold surface and the change of the flow direction although this is constant across all the servers.

The results of simulation show that the server at the top of the rack (`eng01`) receives the coolant with a flow rate of 28% higher than the server in the bottom of the rack (`eng30`) when the CHx40 pumps at the highest speed (5). This variation reduces to about 24% at the lowest CHx40 pumps speed (1). The results also show that the differences between pump speed 4 and 5 are minimal. Since the total flow affects the heat transfer characteristics, with larger flows having a lower thermal resistance and a greater thermal capacity then this corresponds with servers at the top of the rack having a lower CPU temperature than those towards the bottom of the rack.
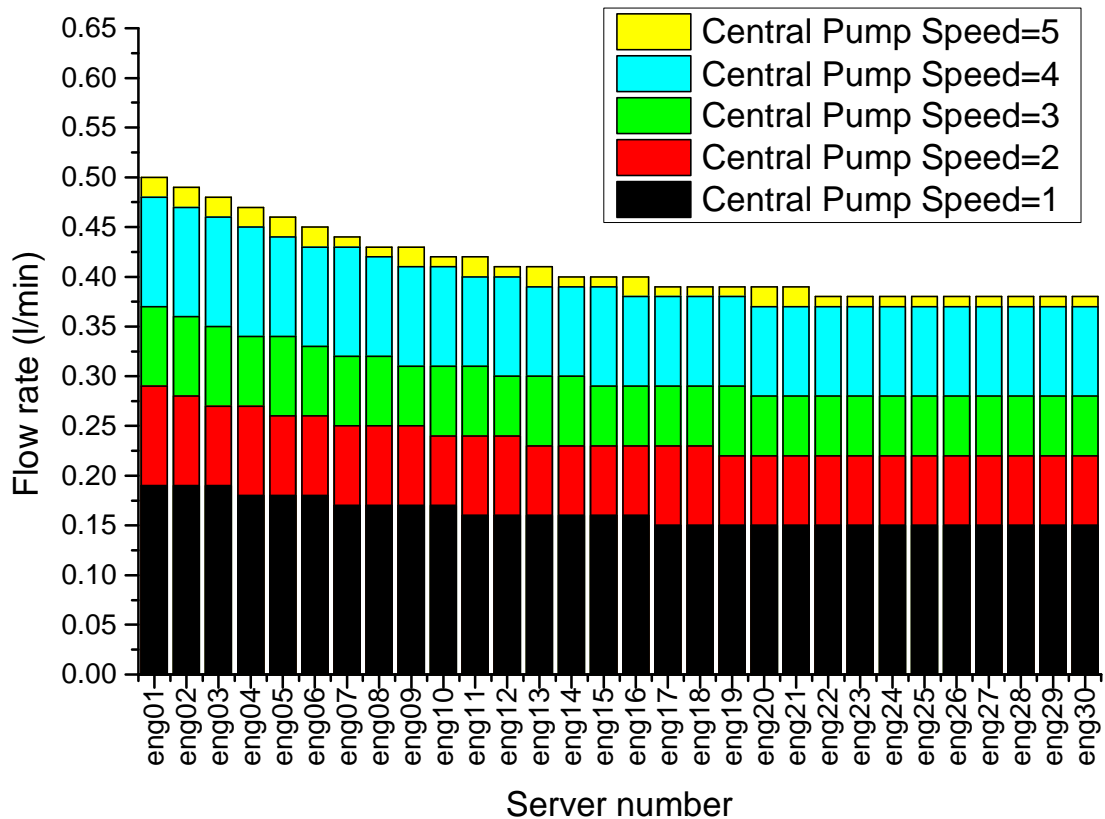
Figure 6.3 Rack flow distribution in the centralised pumping case.

6.2.3 Distributed pumping

The flow distribution in the distributed pumping configuration is shown in figure 6.4 for the case where the CHx40 pumps are deactivated (pump speed 0) and for the cases where the CHx40 pumps are enabled at the five different pump speeds. The flow maldistribution in the distributed pumping configuration shows to be higher than the centralised pumping case. The received flow rate by the server at the top of the rack is higher by 33% than the lowest in the rack. The differences decrease with reducing CHx40 pump speeds with a minimum of 24% when the CHx40 pumps are turned off.

The flow rates through the bottommost ten servers in the rack are almost the same (either for centralised or distributed pumping configurations). It follows that, the CPUs temperature of the lower ten servers in the bottom of the rack should be the same as the flow rate across these 10 servers. The CPU temperature distribution will be discussed in the next section.
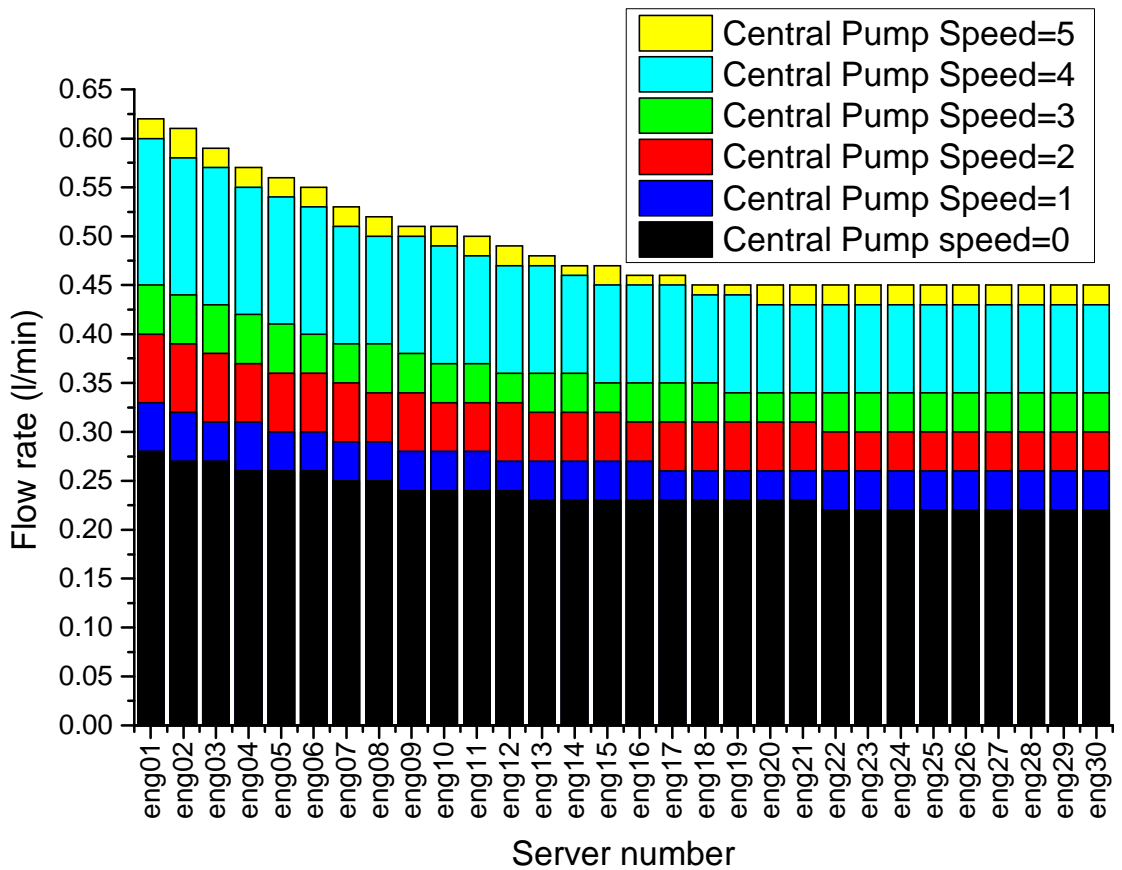
Figure 6.4 Rack flow distribution in the distributed pumping case.

6.2.4 CPU temperature variation

The average CPU temperature distribution of all the servers at 100% utilisation is shown in figure 6.5 with the central pumps speed of 5. The general trends of the servers' CPU temperatures statistically show that the servers at the top of the rack have a lower temperature than the servers in the bottom of the rack. This is proven based on the statistical analysis of the data which show that the temperature distribution can be represented by a quadratic model with a regression factor of 0.72.

The lower temperature is shown to be in the CPUs of server 3 (eng03) of about 42.5℃ and the maximum temperature is shown in the CPUs of server 30 (eng30) which is about 55℃. However, the average temperature of all the CPUs of the rack is about 49℃.
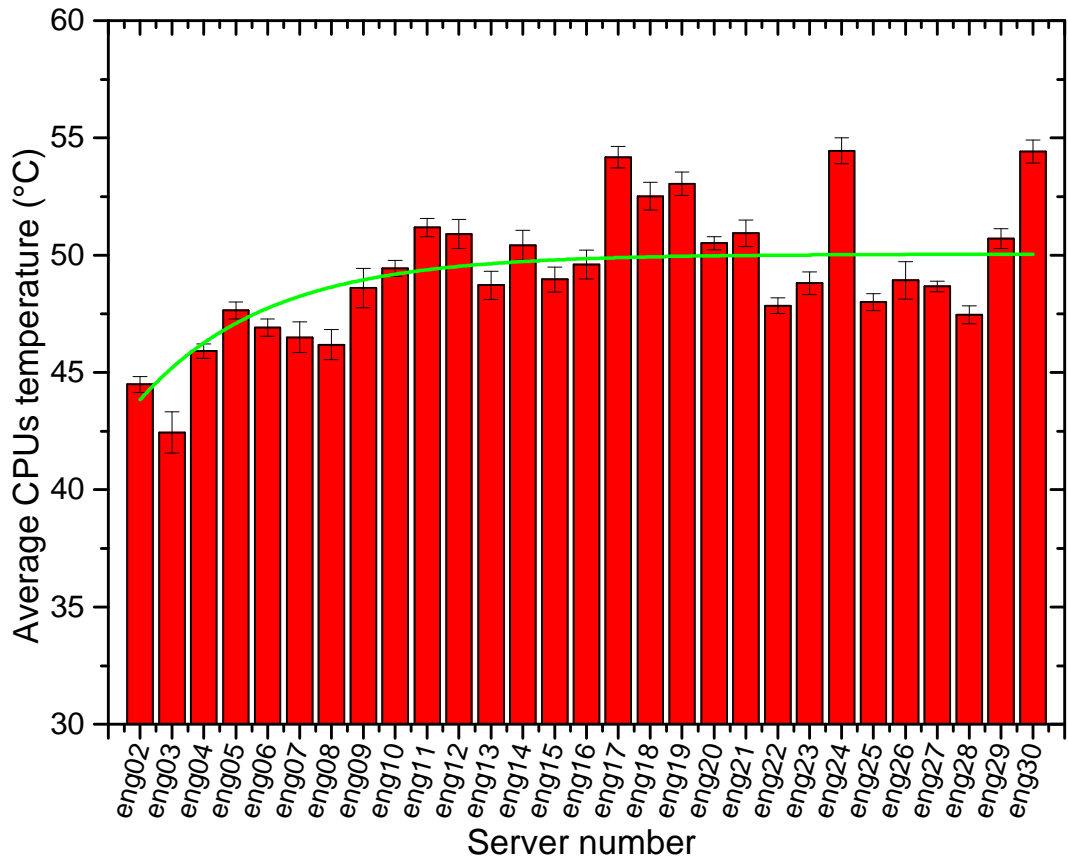
Figure 6.5 CPU Temperature distribution of the servers at 100% load operation.

The temperature distribution of the servers for all the rack at the idle operation is shown in figure 6.6. The CPU temperatures shows high variation between the servers. It is also difficult to show the effect of flow maldistribution as the CPUs are in idle mode and the effect of flow distribution is insignificant and overweighted by other factors such as the pressure drop variation. The minimum temperature is also shown to be in the CPUs of server 3 (eng03) of about 33℃ and the maximum temperature is shown in server 24 (eng24) in the bottom of the rack of about 41℃ and the average CPUs temperature of all the rack is about 37℃.
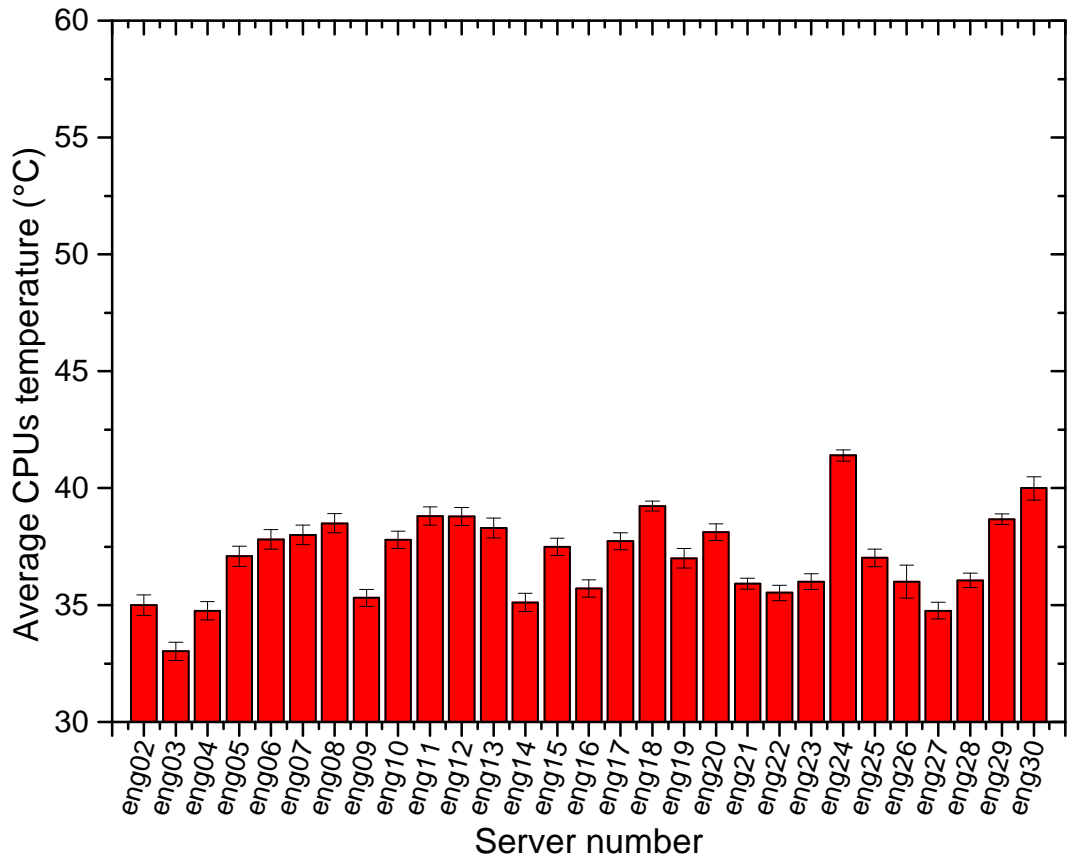
Figure 6.6 CPU Temperature distribution of the servers in idle operation

6.2.5 Variation of average CPU temperatures for the same server flow rates

The results of temperature variations under the same load in the previous section especially in the last ten servers of the rack, which theoretically have the same flow rate, indicate that there are other factors that is also affecting on the flow maldistribution and in turn different CPU temperatures. However, to calculate the effect of the flow distribution on the CPU temperature variation, the experimental CPU temperatures are shown against the theoretical server module flow rates for both cases of full load and idle operation of the servers.

The CPU temperature as a function of server flow rate is shown in figure 6.7 for the 100% utilisation load. It can be seen that the servers of lower flow rate, which are the servers in the bottom of the rack, have a temperature variation from about 47 to 54°C even if the flow rate is the same. This indicates that the

server module also drives the amount of the coolant flow rate that passes through it. This issue is attributed to different factors such as the placement of the server module over the CPU, corrosion and blockage that may occur in the microchannel of the heat sink of the servers' modules. This is supported by a previous recent study by Alkharabsheh et al. [92] where they found that the servers' modules have different flow resistance and that the position of the corrugated tubes also have a high effect on the pressure drop of the servers' modules.

The relationship between the CPU temperatures and the server module flow rates can be best fitted with the polynomial fit of the second order with a regression factor of 0.47. The low regression factor is also attributed to the aforementioned factors that results in different CPU temperatures for the same flow rate through the DCLC heads of different servers.

The CPU temperature as a function of the calculated flow rate for the idle operation of the servers is shown in figure 6.8. The decrease in the CPU temperatures as a function of increasing the server flow rate is within the CPU temperature variation for the same flow rate. These results indicate that the effect of the flow rate variation is statistically within the effect of the variation in server modules pressure drop. Given this, it is difficult to evaluate the effect of flow variation in the idle case as the CPUs of the servers have relatively lower temperature even if the general trend of the CPU temperatures is still decreasing with increasing server flow rates.
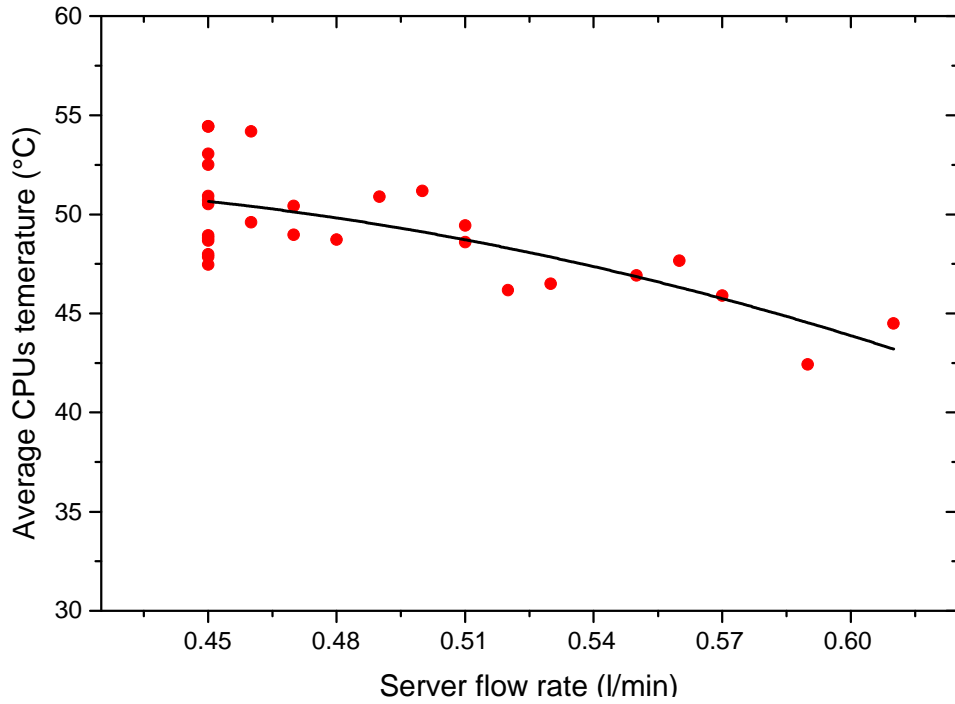
Figure 6.7 Average CPU temperatures as a function of the server flow rates at 100% stress level of the rack.
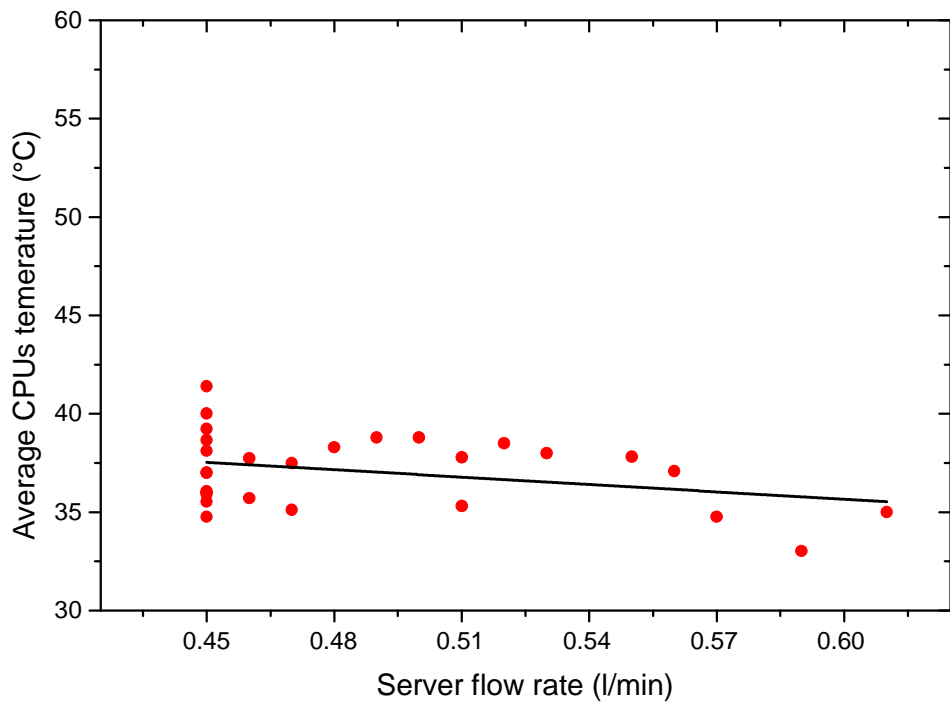
Figure 6.8 Average CPU temperatures as a function of the server flow rates at idle operation of the rack.

6.3 Optimisation of the secondary loop in the rack for uniform flow distribution

One of the requirements in the process of optimising the DCLC servers is to achieve uniform flow distribution through the servers of the rack. Improving the pressure drops in the components of the DCLC servers loop should be included in the optimisation of these systems to achieve uniform flow distribution, reduce the energy consumption and increase the IT performance.

Optimising the DCLC design should take two paths: Improving the servers' module design to reduce the pressure drop in these segments of the loop and improving the design of the supply and collection manifolds. The scope of the current study will focus on optimising the manifolds design to reduce the pressure drop to achieve uniform flow distribution.

Figure 6.9 shows the resulted improvement between the flow distribution and the equivalent manifold diameter as well as the improvement in the total loop flow rate for the central pumping of the system. It can be seen that the percentage difference between flow rates of the server in the top of the rack and the server in the bottom of the rack diminishes from 28% to 2% when the equivalent hydraulic diameter of the manifold is doubled from 2.5cm to 5cm. For this, the total loop flow rate improves by about 9%, as well as reducing the maldistribution of the flow.

Figure 6.10 shows the decrease in the percentage difference between the flow rate of the coolant through the server module in the top of the rack and the server module in the bottom of the rack for the distributed pumping configuration. The total improvement of the flow rate is also presented for increasing hydraulic diameter of the manifolds. It is shown that the percentage difference between the highest and lowest flow rates through the servers' modules improves with increasing the hydraulic diameter. The total loop flow rate improves by about 10% when changing the hydraulic manifold diameter from 2.5 to 5cm.

The results of optimising the cooling loop are crucial when designing the overall system to improve the energy efficiency and increase the reliability and

robustness of using the DCLC servers in the data centre industry. The results indicate that doubling the hydraulic diameter from 2.5cm to 5cm will have a high impact in achieving uniform flow for all the servers which are connected in parallel in the loop.
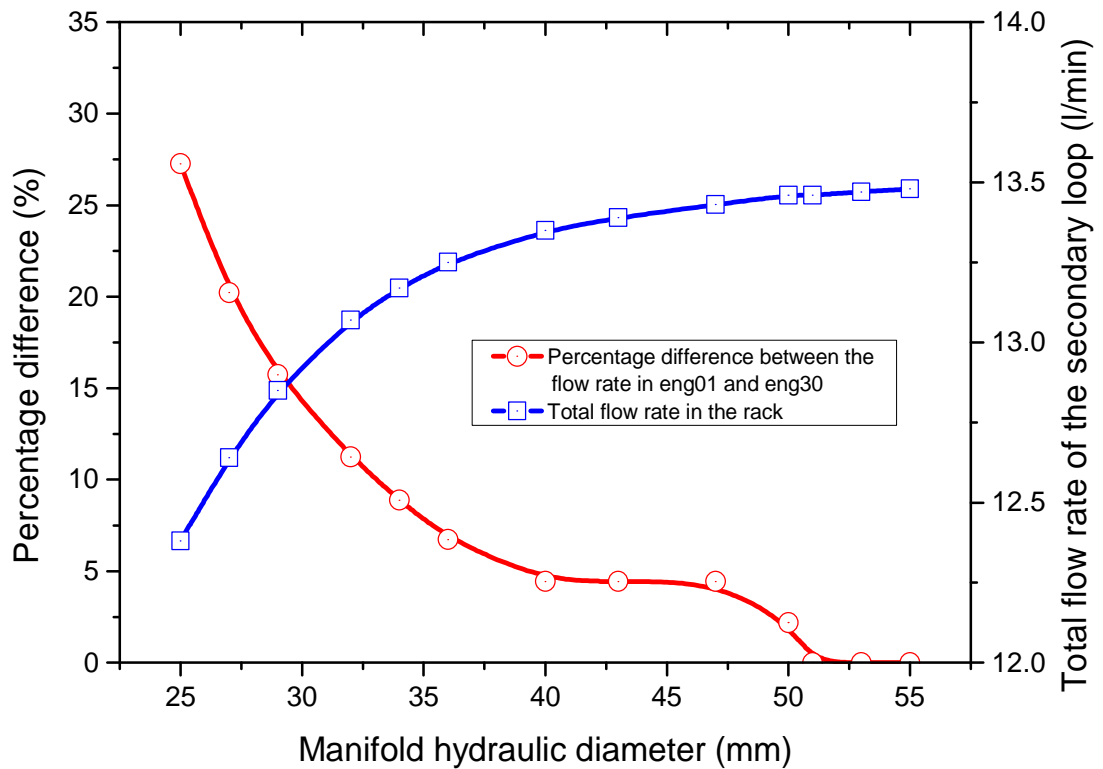
Figure 6.9 Manifold size effect on the flow distribution in the central pumping configuration (the CHx40 pumps are running only).
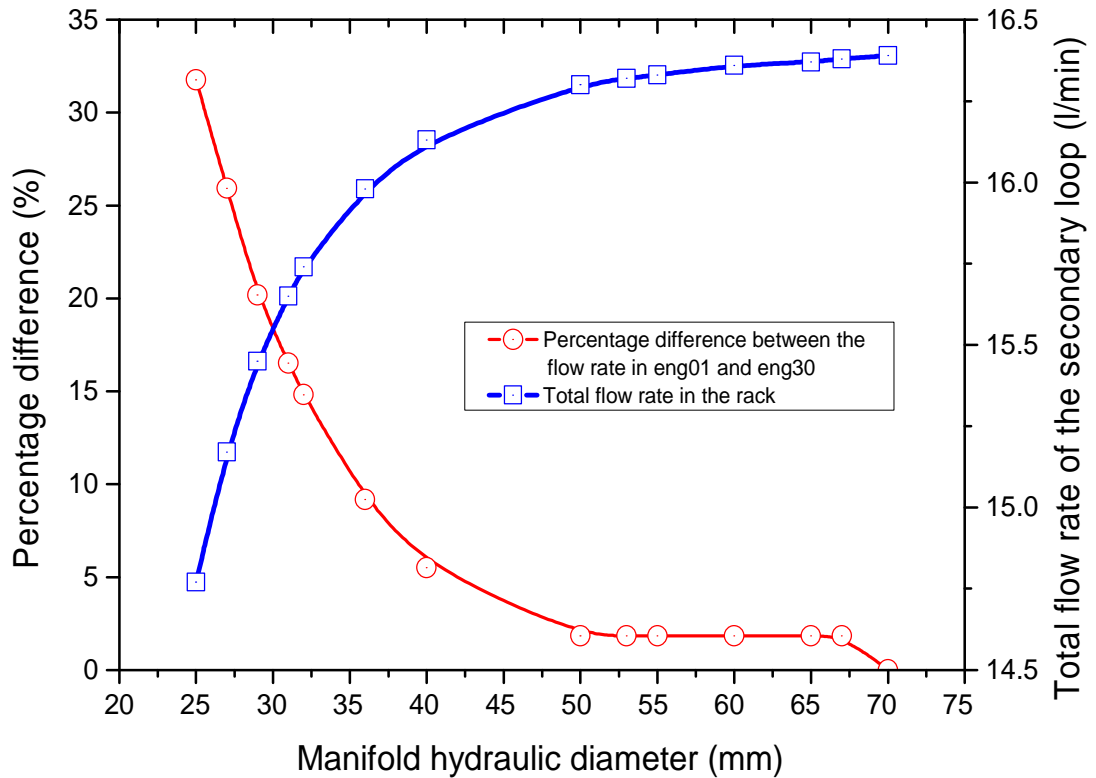
Figure 6.10 Manifold size effect on the flow distribution in the distributed pumping configuration (the CHx40 pumps are kept running as well as the small pumps at the CPUs).

6.4 CPU temperatures under centralised and distributed pumping

This section presents experimental comparison results of the effect of the pumping configuration and the flow rate on the CPU temperatures. As mentioned in chapter 4, four servers are selected to represent server CPU temperatures through the experiments which are `eng02`, `eng10`, `eng20` and `eng30`. The rack inlet temperature was set through all the experiments which is regulated by the AHU. The flow rate in the secondary loop was regulated from the CHx40 pumps. The experiments were carried out for five different load levels which are idle, 25%U, 50%U, 75%U, and 100%U.
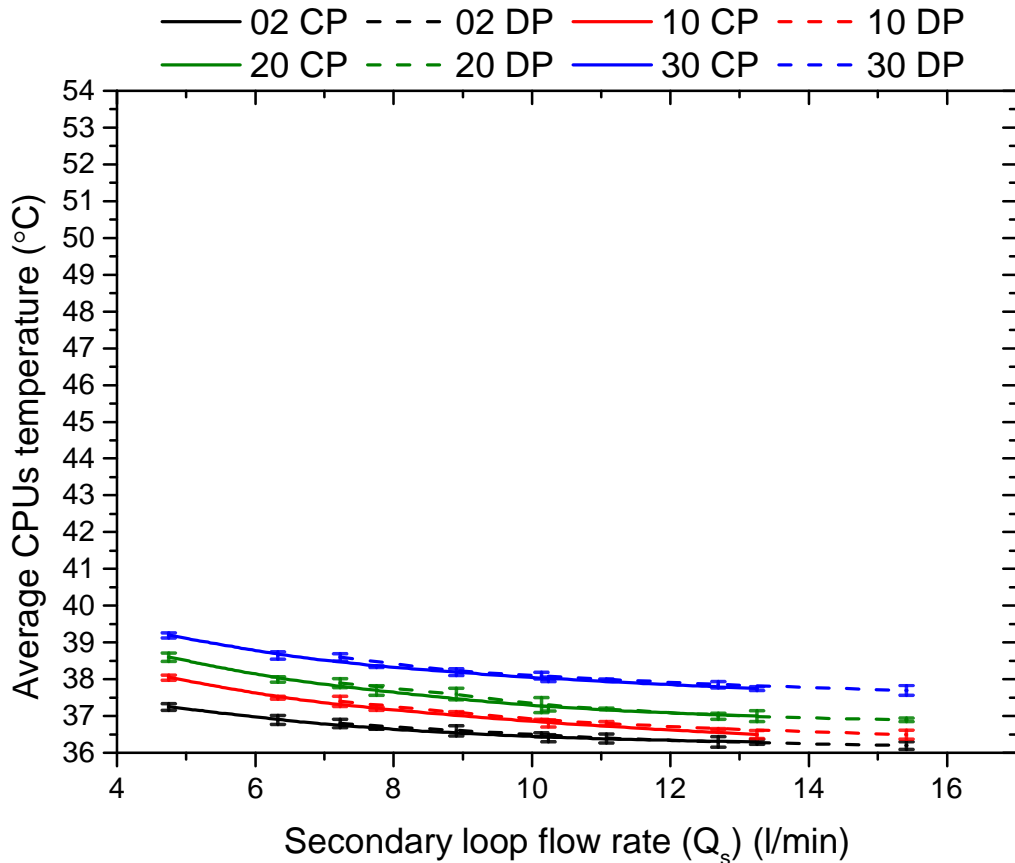
Figure 6.11 Average temperature of selected servers CPUs as a function of secondary loop flow rate under idle operation (the numbers in the graph legend are referred to the servers number from top to bottom of the rack while the CP and DP refer to the centralised and distributed pumping respectively).

As shown in figures 6.11 to 6.15, the CPU temperatures on average are higher for the distributed pumping (DP) than the central pumping (CP) for the same flow rate. The differences between the DP and the CP configurations is within the percentage error for the four selected servers in the idle case as shown in figure 6.11. The improvement in the CPU temperatures with increasing the coolant flow rate is relatively small and within 1℃ for both the CP and DP configurations.
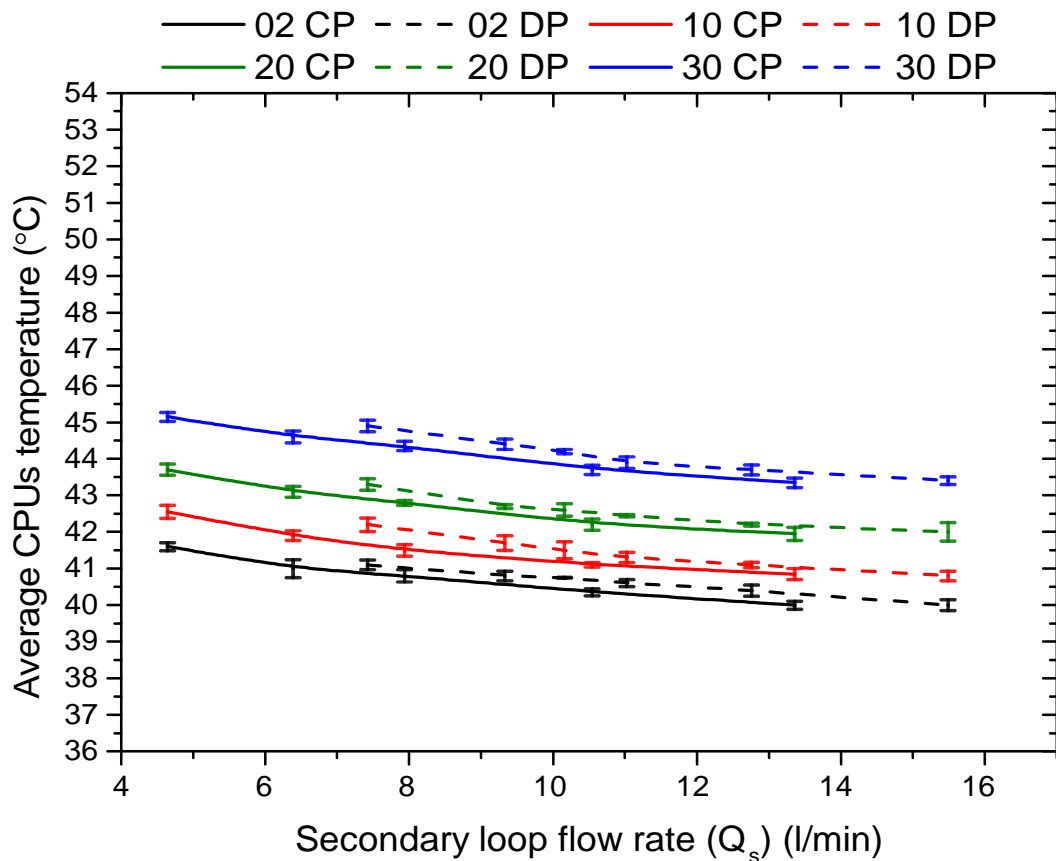
Figure 6.12 Average temperature of selected servers CPUs as a function of secondary loop flow rate at 25% utilisation (the numbers in the graph legend are referred to the servers number from top to bottom of the rack while the CP and DP refer to the centralised and distributed pumping respectively).

The differences between the CPUs temperature of the CP and DP is shown to be greater when the servers are loaded to 25% stress level as compared to idle, as shown in figure 6.12. However, the difference between the DP and CP is still minimal, at ~1% and also within the error bars of the experiment. The improvement in the CPU temperatures was found to be about 4.2% on average for increasing the flow rate from the minimum value of about 4.7 l/min to about 13.5 l/min for CP case and of about 4% for increasing the flow rate from 7.2 l/min to about 15.8 for DP case.

At 50% load, the CPU temperatures show a high temperature increase of about 1.5°C on average for the DP than the CP configuration as shown in figure 6.13. The response of the CPU temperatures is higher to the loop flow rate than the idle and 25% load levels.

192
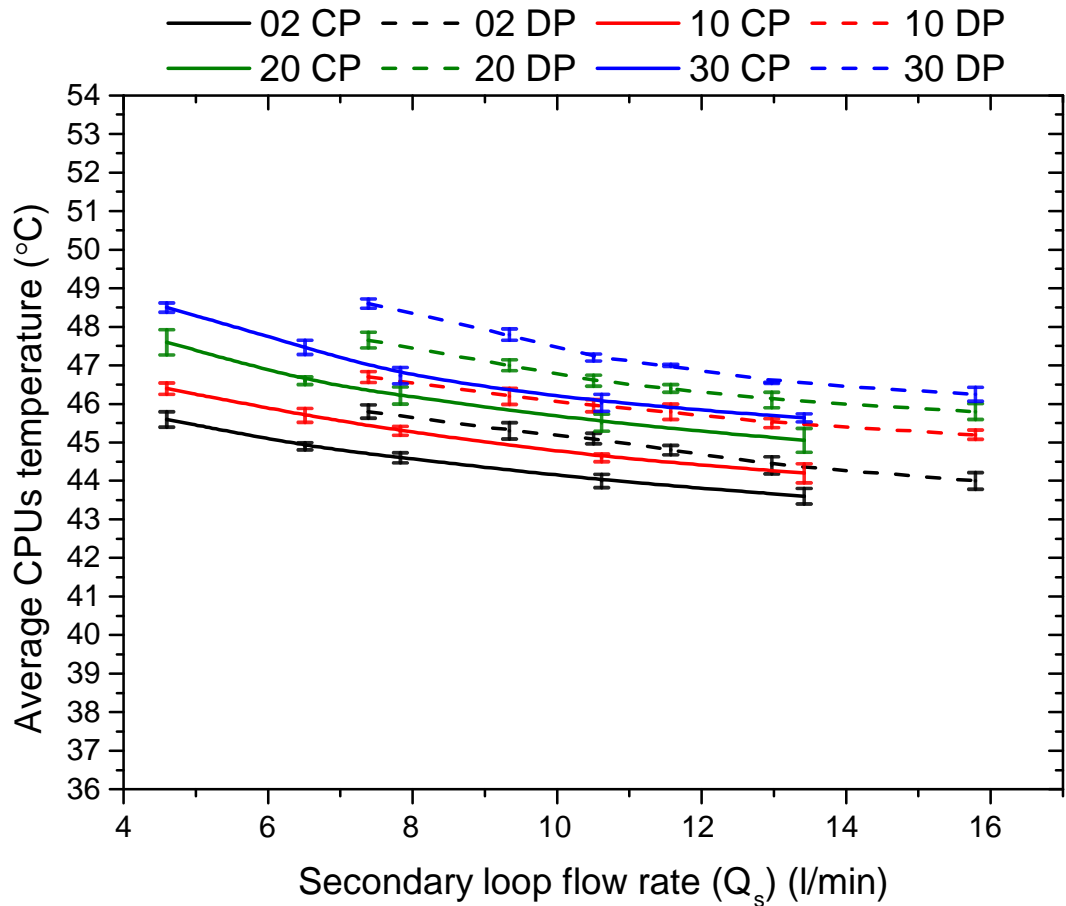
Figure 6.13 Average temperature of selected servers CPUs as a function of secondary loop flow rate at 50% utilisation operation (the numbers in the graph legend are referred to the servers number from top to bottom of the rack while the CP and DP refer to the centralised and distributed pumping respectively).

For the high IT load levels of 75% and 100% utilisation, the differences between the CPU temperatures in the CP and DP is more evident and as high as 2°C as shown in figures 6.14 and 6.15, respectively. The average reduction in the CPUs temperature is of about 6% when increasing the pump speed from 4.3 to 13.7 l/min for the CP configuration and of about 4% when increasing the flow rate from the 7.5 to 15.8 l/min for the DP configuration.

It should be mentioned that reducing the coolant flow rate shows a small increase in the CPU temperatures of the selected servers, although two of the servers encountered shutdown during the high load level of operation with the low flow rate in the CP configuration. However, for the DP configuration, all

the servers remained within the safe margin of CPU temperature operation even at the lower flow rate where the CHx40 pumps were deactivated.
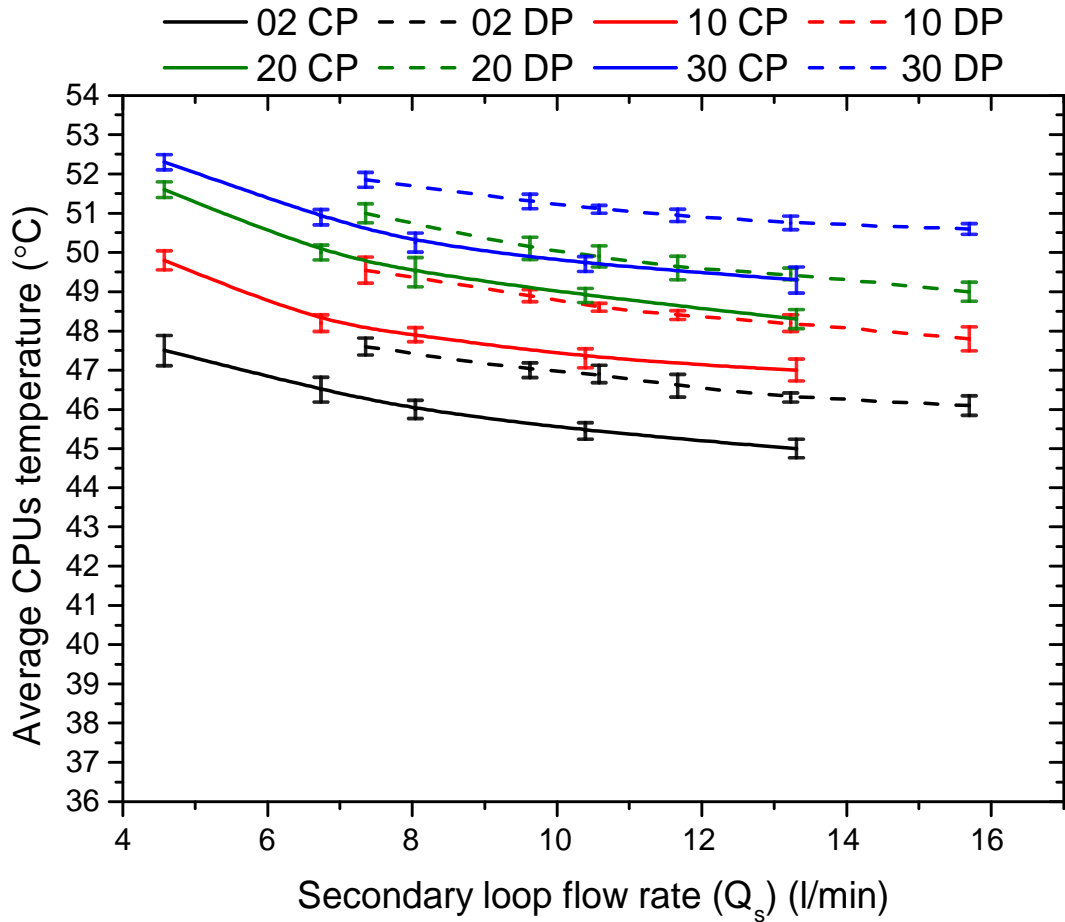
Figure 6.14 Average temperature of selected servers CPUs as a function of secondary loop flow rate at 75% utilisation operation (the numbers in the graph legend are referred to the servers number from top to bottom of the rack while the CP and DP refer to the centralised and distributed pumping respectively).

It can be concluded from the results of this section that activating the DPs should be automated with a specific algorithm to activate the local CPU pumps when the CPU reaches the critical temperature to aid in reducing the pressure drop of the server module and increase the flow rate through that particular server.
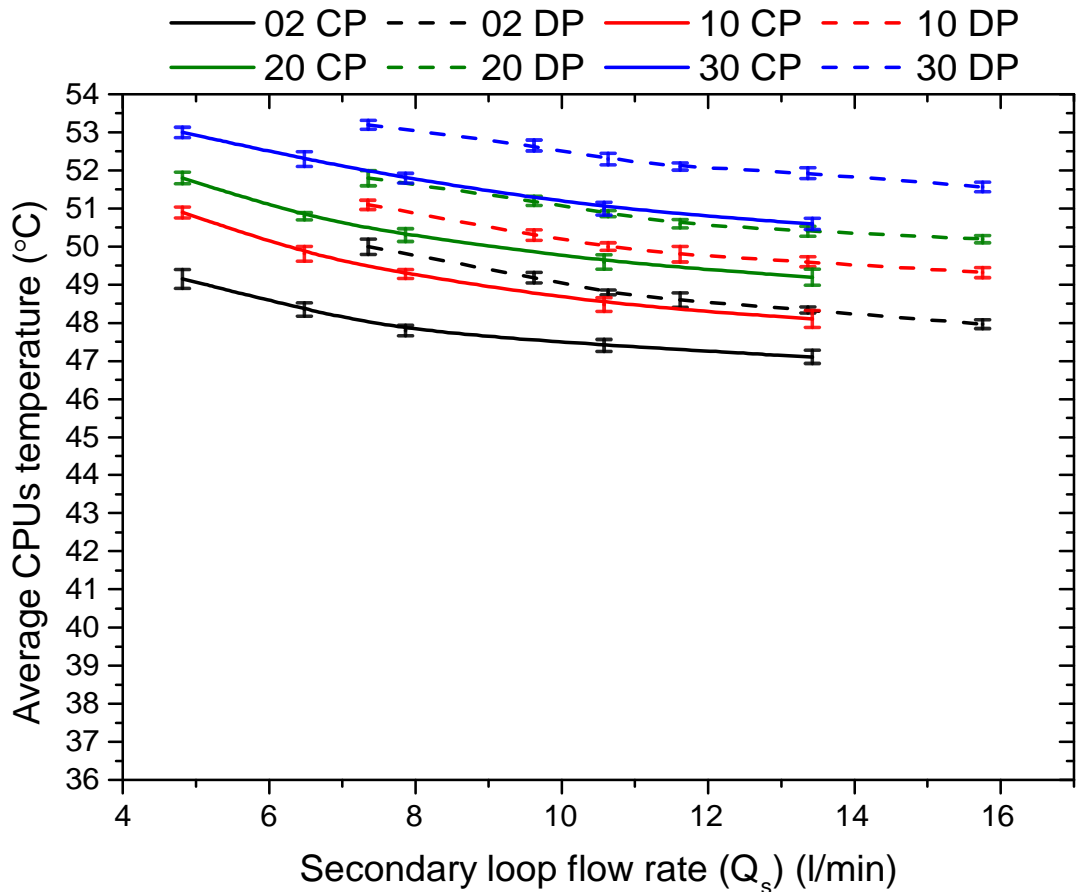
Figure 6.15 Average temperature of selected servers CPUs as a function of secondary loop flow rate at 100% utilisation operation (the numbers in the graph legend are referred to the servers number from top to bottom of the rack while the CP and DP refer to the centralised and distributed pumping respectively).

6.5 Effectiveness of the data centre under centralised and distributed pumping

The effectiveness of the rack is calculated in terms of pPUE and TCUE as explained in chapter 4. The pPUE evaluates the effectiveness based on the total power consumption of the data centre to the power consumed by the rack regardless of the other cooling means inside the servers such as the fans and the server module pumps in cases when they are active. However, the TCUE is also a term of effectiveness, but it takes into considerations all the power consumed by the cooling devices inside the IT. As a result, the TCUE is more

accurate in terms of describing the actual power consumed by cooling of data centre and the power consumed by the IT itself to carry out the operations.

Figure 6.16 shows the pPUE and TCUE in the idle operation of the servers under the two types of coolant pumping: centralised or distributed. The pPUE is lower in the DP case than the CP case for the same range of the flow rate of the secondary loop. The reason is that in the DP configuration, the servers' modules pumps are activated which are assisting the CHx40 pumps and also reduces the pressure drop in the servers' modules. Moreover, from the pPUE definition, the power consumption by the servers' modules are added to the IT side which in turn reduces the calculated power consumed by the cooling side and as a result lower pPUE. In the TCUE calculations, the case is different from the pPUE where the power consumed by the fans and the servers' modules pumps are added to the cooling side. As a result, the TCUE deteriorate in the case of the DP configuration than the CP configuration for the same IT load as shown in figure 6.16.

The effectiveness metrics for the full operational load of the IT is shown in figure 6.17 for the centralised and distributed pumping configurations. The same behaviour is observed of an improvement in the pPUE metric when activating the server module pumps while TCUE deteriorates. However, there is a slightly higher increase in the TCUE of DP and CP in the full load than the idle case. This is due to the increase in the fan speeds of the servers which casuses deterioration to the TCUE by increasing the power consumed by cooling inside the servers. However, this increases in the power consmuption of the servers is due to the increase in the fan speeds improving the pPUE as it is considered as part of the IT side in this case. It should also be noted that the power consumed by the server modules is a fixed value regardless of the CPU temperatures.
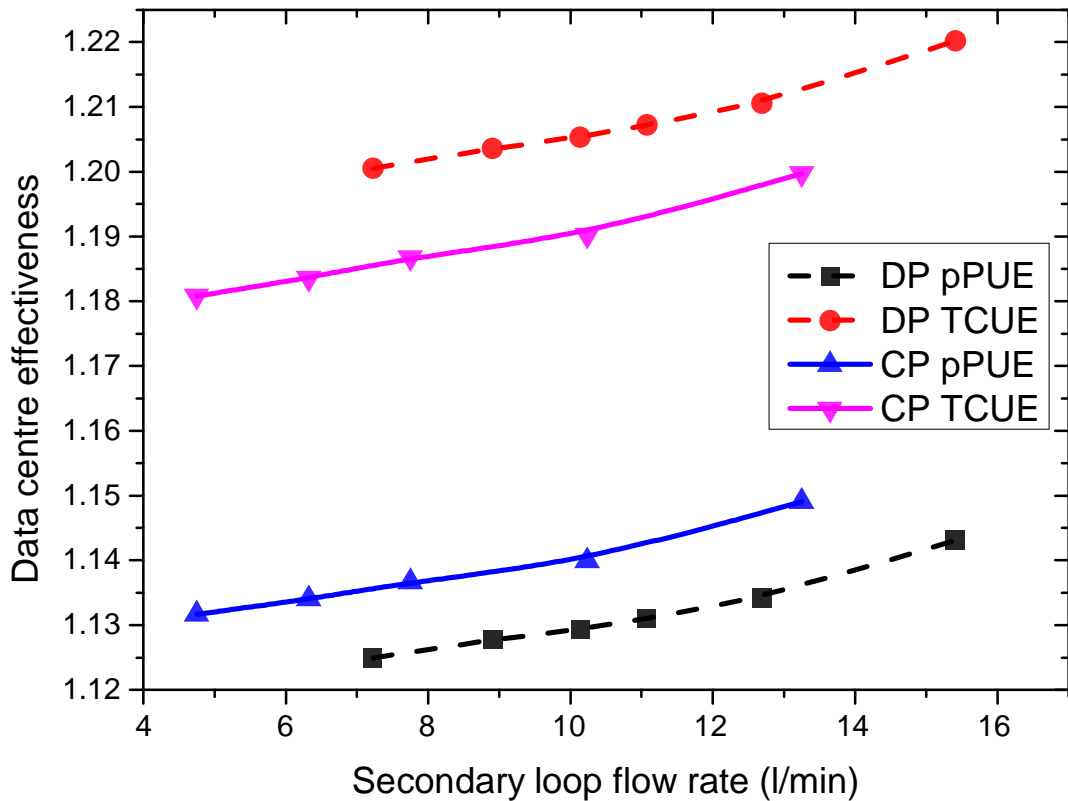
Figure 6.16 pPUE and TCUE of the rack level as a function of the secondary loop flow rate for the idle operation under centralised and distributed pumping (the CP and DP refer to the centralised and distributed pumping respectively).

It can be seen from figures 6.16 and 6.17 that both effectiveness metrics increase with increasing of the secondary loop flow rate for both pumping configurations of distributed and central. The reason for this is that increasing the secondary flow rate increases the power consumed by the CHx40 which is always part of the cooling side of the effectiveness metrics. It should be mentioned that increasing the secondary loop flow rate have an improvement in CPU temperatures of about 2℃ as noticed in the previous section for the 100%U load. The temperature reduction in the CPU temperatures is accompanied by a decrease in the power consumption of the server as will be investigated in detail in Chapter 7. However, the improvement of increasing the flow rate of the secondary loop to reduce the power consumption is insignificant.

Providing sufficient flow rate to the servers is necessary for safe and reliable operation. The experiments revealed that reducing the coolant flow rate in the rack less than 8 l/min could lead to rising problems in some of the servers that have high pressure drop in their cold plate modules. This issue was observed in the high IT load experiments although the tests that were carried out in the cooler temperature margin of ASHRAE W4 [20] recommendation for liquid cooling.
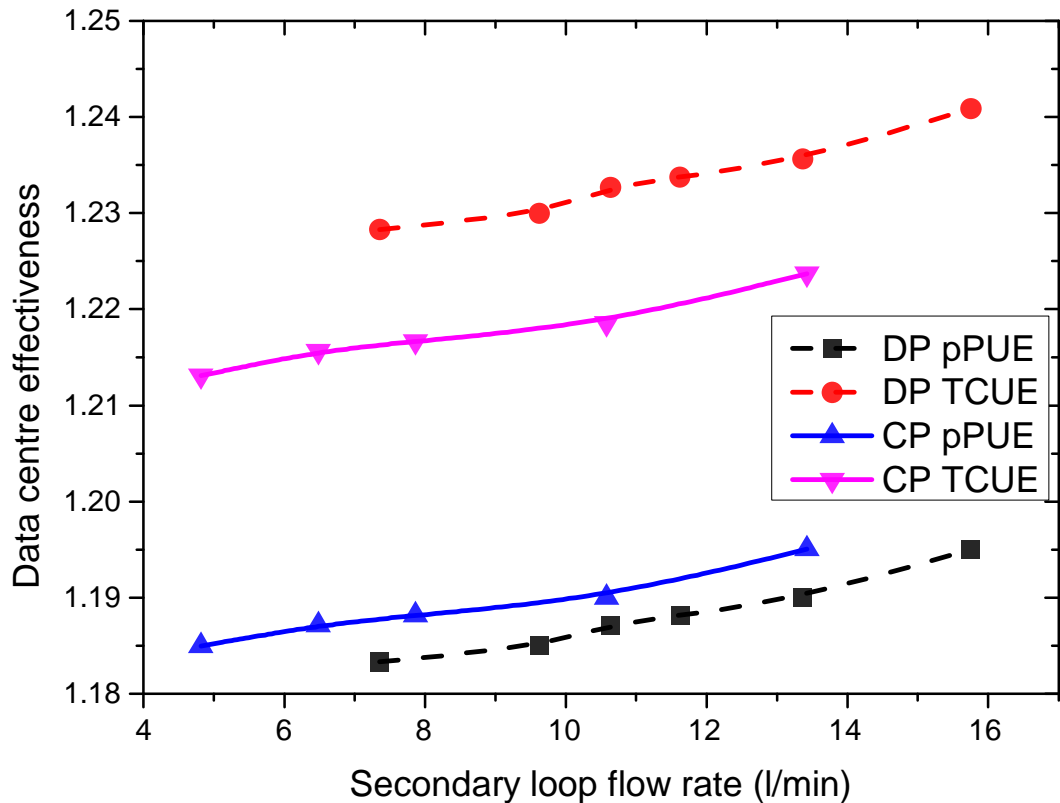
Figure 6.17 pPUE and TCUE of the rack level as a function of the secondary loop flow rate for the 100% IT load under centralised and distributed pumping (the CP and DP refer to the centralised and distributed pumping respectively).

6.6 Summary

This chapter aimed to address the issue of flow maldistribution and the relationship with the components temperature variations. This was studied to optimise the design to obtain uniform flow for all the servers. The coolant pumping configuration and the flow rate in the rack was also investigated for various possible IT loads in terms of the component temperatures. The pPUE and TCUE metrics of the data centre were calculated for each case to produce a general understanding of the effectiveness of the design and power consumption.

The EPANET open source software was used to analyse the flow in the rack for the two types of pumping configurations: centralised and distributing pumping. The model predicted the flow rate in the rack with a good accuracy of 2.3% compared with the experimental results. The results of the analysis showed that the server in the top of the rack (eng01) receives higher flow rate than the server in the bottom of the rack (eng30) and this difference is 28% in the centralised pumping and 30% in the distributed pumping case. The differences in the coolant flow rate received by every server resulted in a general increase in the CPU temperatures of the servers from the top to the bottom of the rack.

The flow maldistribution and the resulting temperature differences of the CPUs of different server positions in the rack is attributed to the friction losses in the manifolds and different pressure drop of different DCLC heads. Optimisation analysis on the discharge and collection manifolds showed that increasing the hydraulic diameter from 2.5 to 5cm provides uniform flow for all the servers and increases the total rack flow rate of about 10%.

The CPU temperatures in the rack was found to be higher for the distributed pumping configuration than the central pumping configuration at the same flow rate. This increase becomes more evident at the high IT loads where it showed a difference of 2℃ on average for the 100% utilisation of the rack. Increasing the total flow rate of the rack showed less effect on the components temperature at the low level loads while it was more significant in the high

level loads. However, two servers encountered shut down when the rack flow rate was reduced to a minimum at the 100% utilisation load in the central pumping configuration case. Therefore, the server pumps can be controlled using an algorithm in future designs to activate the module server pump when the CPU temperature reaches a specific critical value to increase the flow rate for that server.

The pPUE and TCUE metrics increase with higher flow rate as more power is added to the cooling side by the CHx40. The TCUE is more sensitive to the IT load as in the high loads the fan speeds ramped up which adds more power to the cooling side and hence higher TCUE whilst for the pPUE, the server fans are added to the IT side and hence the pPUE decreases.

Chapter 7 Impact of the rack inlet temperature on the DCLC servers performance

7.1 Introduction

In this chapter, the 30 Sun Fire V20z servers are tested at a rack level to investigate the effect of the inlet water temperature to the rack on the CPU temperatures, rack power consumption, IT performance, effectiveness, and the data centre workload power efficiency. The effect of the inlet water temperature will also be tested for a range of IT loads (from idle to full). This study will also serve to optimise the DCLC servers in terms of weighting between the power consumption of the data centre and the computational efficiency of the servers.

The general trend of the data centre industry is to reduce the power consumption of the cooling facilities by increasing the supplied temperature to the IT thereby requiring less work within the cooling facility to provide the targeted temperature. This will result in a lower power consumption and higher effectiveness of the data centre. However, higher inlet water temperature to the servers increases the operating temperature of the CPUs which can have a negative impact on the power consumption of the servers and the operational performance of the CPUs [150].

The main goal of this chapter is to critically evaluate the performance of the rack under a range of operating temperatures within the ASHRAE W4 envolope [151] for data centre liquid cooling. This will also include examining the power consumption required for cooling, the power consumption of the rack and the computational performance of the servers. Following these objectives, the effect of the rack inlet temperature on the CPU temperatures of selected servers will be presented first. Then the effect of the rack inlet temperature on the rack power consumption and computational performance will be presented. Then the effectiveness of the data centre in terms of pPUE and TCUE will be given. Finally, the pDWPE which shows the compromise

between the effectiveness and the efficiency of the data centre will be discussed.

7.2 Brief review of the experimental methodology

As described in chapter 4, the AHU is set on auto mode to adjust the inlet temperature to the rack at the required setpoint. The temperature range is between 15℃ and 40℃ which is the temperature delivered to the RDHx before passing into the CHx40. For this range of temperatures, the inlet temperature to the rack from the CHx40 is within the ASHRAE W4 recommended range for server liquid cooling. CPU temperatures of selected servers are chosen to be presented. Power measurements will also be taken for all the components of the system from the servers to the AHU. The computational operation of a selected server will be calculated per unit of power consumption of the selected server. The results of experiments in this chapter are for distributed pumping configuration.

7.3 Idle operation of the rack

Five servers are selected to present the average CPU temperatures for the tested range of rack inlet coolant temperature. These servers are `eng02`, `eng06`, `eng10`, `eng20` and `eng30`. These servers are selected to present the different servers in different levels of the rack from top to bottom.

The selected servers CPU temperatures show a clear increase with increasing the rack inlet temperature as shown in figure 7.1. The average CPU temperature is about 34℃ when the inlet water temperature to the rack is 20℃ and increases to about 52℃ when increasing the rack inlet water temperature to 38℃. The relationship between the CPUs temperature and the rack inlet coolant temperature is shown to be linear. Again the variation between the CPU temperatures of different servers is due to the flow maldistribution inside the rack that resulted from the pressure drop in the manifolds and the variation of the pressure drop in the different servers' modules as mentioned in chapter

202

6. It should also be mentioned that all the selected servers respond to all the inlet water temperatures consistently.

Figure 7.1 Average CPUs temperature of the five selected servers in the rack as a function of the rack inlet water temperature in the idle operation of the servers.

In the idle operation of the servers, the CPUs are not performing any computational calculations hence the Flops per watt is zero for all the experiments regardless of the servers supplied temperature as shown in figure 7.2. The average power consumption of the rack is measured over the experiments for different inlet water temperature as shown in figure 7.2. The increase in the total power consumption of the rack when the inlet water temperature was stepped from 20℃ to 38℃ is very marginal of less than 0.5% (about 19W). The reason for this is that the servers are not under any computational process and the slight increase is only due to the slight increase in the static power in keeping the server on.

The effectiveness metrics of the pPUE and TCUE were also calculated as shown in figure 7.3. The pPUE and TCUE are very high when the temperature

inlet to the IT is set to 15℃. The reason for this is that these experiments were carried out in summer ambient conditions which were challenging for the AHU to reach the low target temperature even with the servers are in idle operation. Hence the power consumption by the AHU was high whereas the power consumed by the rack was relatively small which results in a high pPUE. The TCUE follows the same trend as the pPUE with a fixed difference between them for all the range of temperatures as the servers fans were kept relatively constant as the temperature remained within the targeted margin of the operation of the servers.

It can also be seen from figure 7.3 that the pPUE and the TCUE are decreasing with increasing the rack inlet temperature. This is due to the decrease in the power consumption of the AHU as the required work to achieve the targeted temperature is reduced and the increase in the power consumption of the rack is minimal which have almost no effect on the power effectiveness metrics.

Figure 7.2 Workload power efficiency and rack power consumption as a function of the rack inlet temperature in the idle operation.

Figure 7.3 Effectiveness of the rack level as a function of the rack inlet water temperature for the idle operation.

7.4 Moderate IT load of the rack

For the moderate IT load experiments, the rack is tested under two stress levels: 25% utilisation and 50% utilisation. In both experiments, the increase in the rack inlet temperature shows a high increase in the CPU operating temperatures. For the 25% utilisation load, the average CPU temperature is about 38℃ when the coolant inlet temperature is 20℃ and increases to about 58℃ when the coolant inlet temperature increases to 41℃ as shown in figure 7.4. On the other hand, the average CPU temperature is about 41℃ when the coolant inlet temperature is 22 ℃ for the 50% utilisation load experiments and increases to become 62℃ when increasing the coolant inlet temperature to 43℃, as shown in figure 7.5. It should be mentioned that server 28 (eng28) was shut down when the inlet temperature to the rack increased to 37℃ and another server (eng27) was also shut down when the inlet rack temperature increased to 41℃ under the 50% utilisation load experiments.
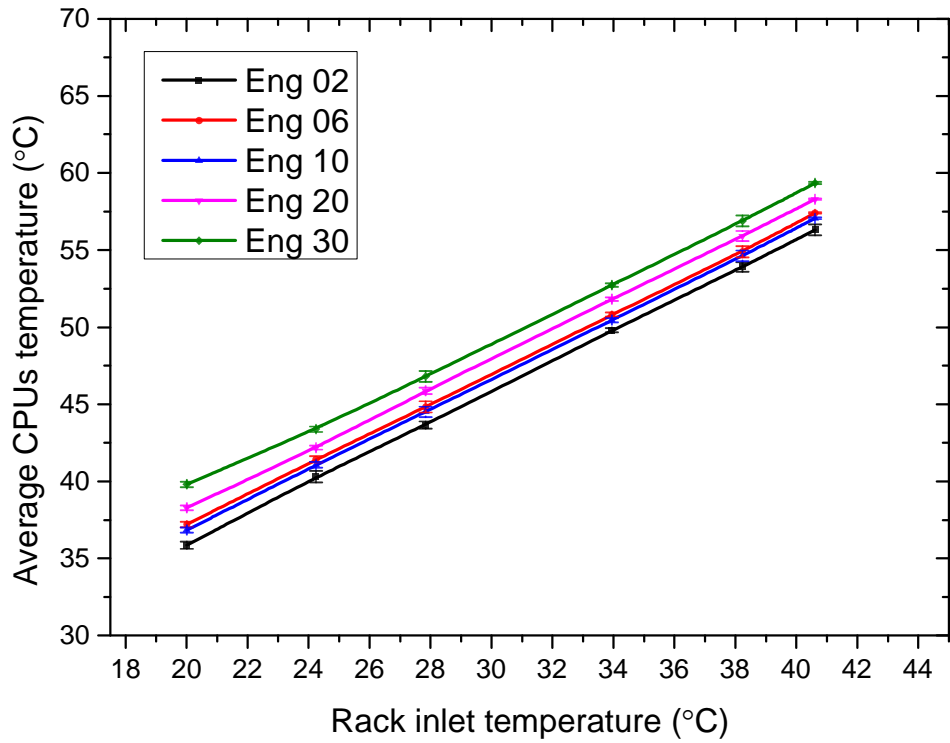
Figure 7.4 Average CPU temperature of the five selected servers in the rack as a function of the rack inlet water temperature at 25% load.
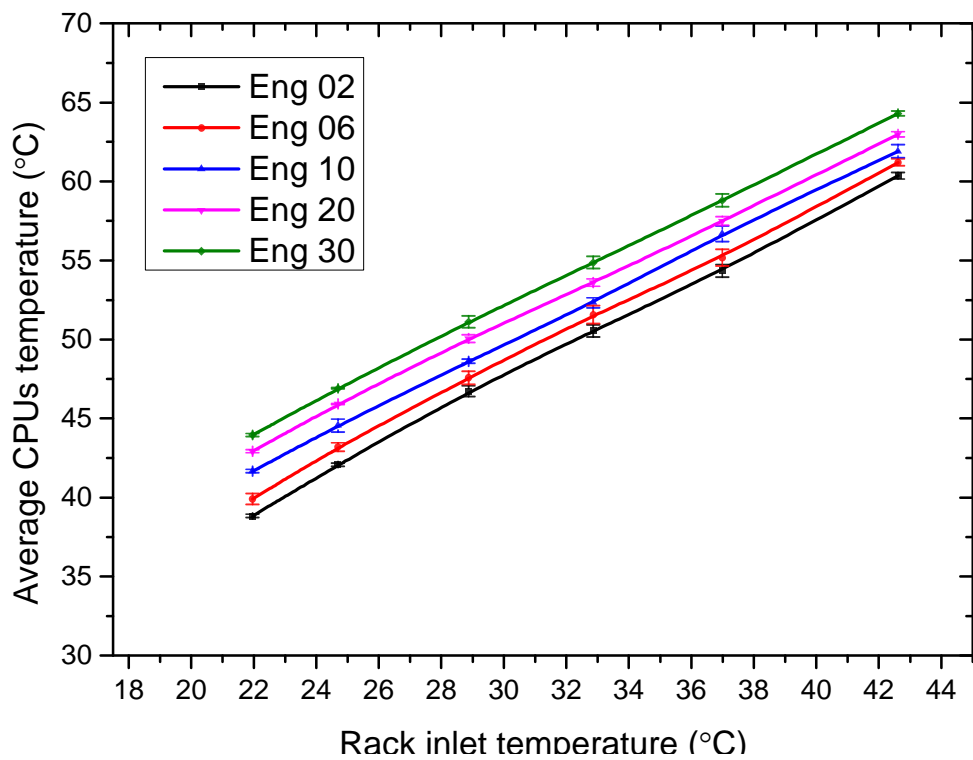
Figure 7.5 Average CPU temperature of the five selected servers in the rack as a function of the rack inlet water temperature at 50% load.

The workload power efficiency and the rack power consumption measurements were calculated for the two moderate stress levels (25% and 50% utilisations) under various rack inlet water temperature. For the 25%U experiments, the WPE was considerably reduced by about 5% when increasing the rack inlet temperature from 20 to 41℃. Whereas, the power consumption of the rack increased by about 5% for this increase in temperature. On the other hand, for the 50%U load, the WPE decreased by about 3% and the rack power consumption increased by about 4% when the rack inlet temperature stepped from 22 to 43℃. The server fan power consumption is included in the rack power consumption in these calculations. Hence the increase in the rack power consumption is due mainly to the increase in what is called the sub-threshold leakage, which is a leakage mechanism of power that is directly related to temperature, and the ramp up of the server fan speeds when the CPU temperatures increase.
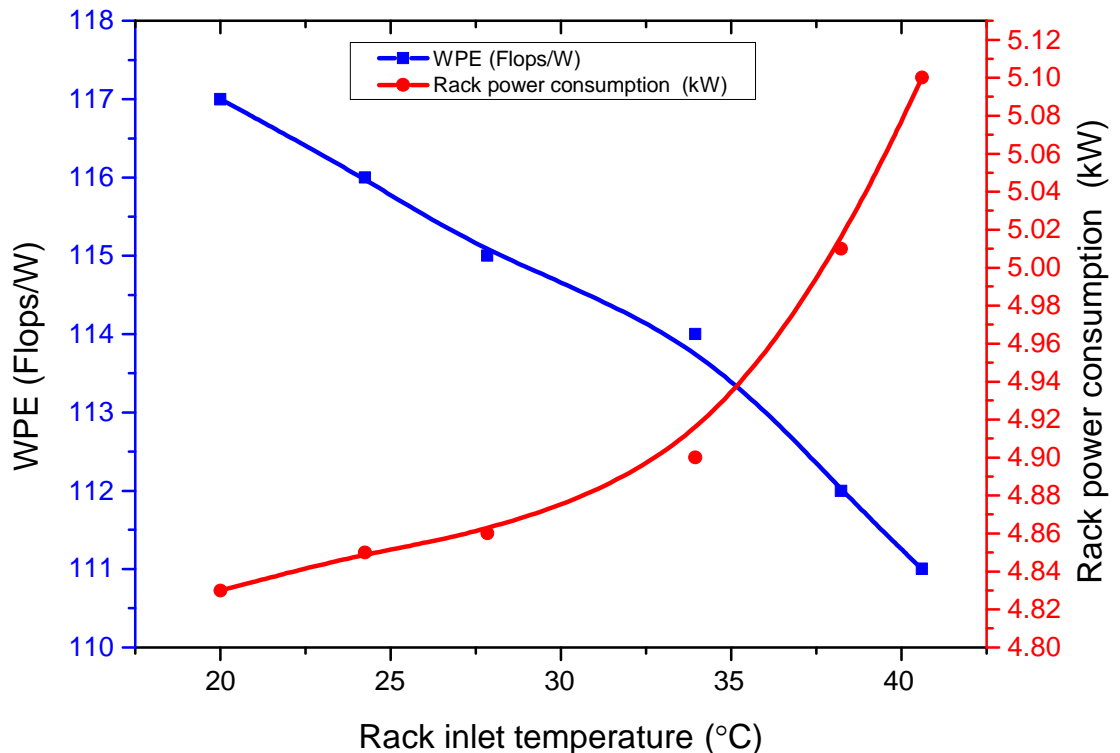
Figure 7.6 workload power efficiency and rack power consumption as a function of the rack inlet temperature at 25% load.

207
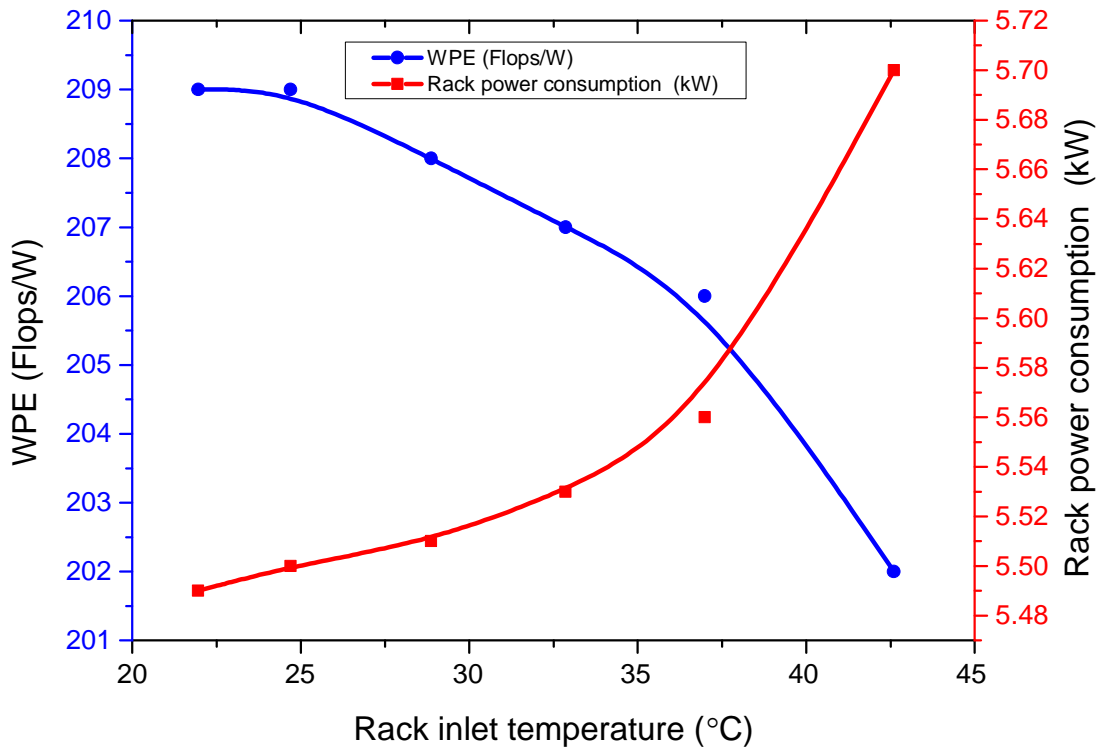
Figure 7.7 workload power efficiency and rack power consumption as a function of the rack inlet temperature at 50% load.

The corresponding effectiveness metrics for the 25% and 50% utilisation experiments are shown in figures 7.8 and 7.9, respectively. The reduction in the AHU power consumption due to the reduction in the work required by increasing the target temperature improves the pPUE and the TCUE. It should be mentioned that the increase in the power consumption of the rack due to the increase of the CPU temperatures as a result of increasing the rack inlet temperature results in an improvement in effectiveness metrics, but it has a negative impact on the total power consumption of the data centre. It was also observed that the speed of the server fans increases when the rack inlet temperature is high even if the DCLC adequately cools the CPUs. These fans are responsible for cooling the RAM that is associated with each CPU. For this, the rate of the improvement of TCUE reduces when the inlet temperature to the rack is higher than 28℃ due to the ramp up of the server fans which adds extra power consumption to the servers, but it will be added to the cooling side when calculating the TCUE.
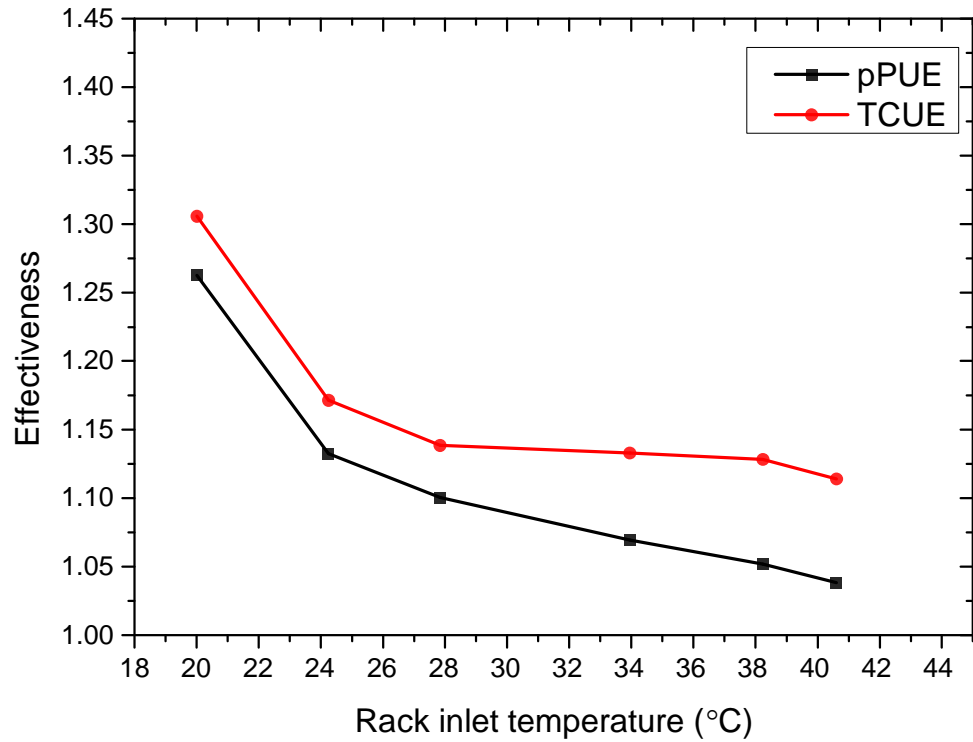
Figure 7.8 Effectiveness of the rack level as a function of the rack inlet water temperature for the 25% load operation.
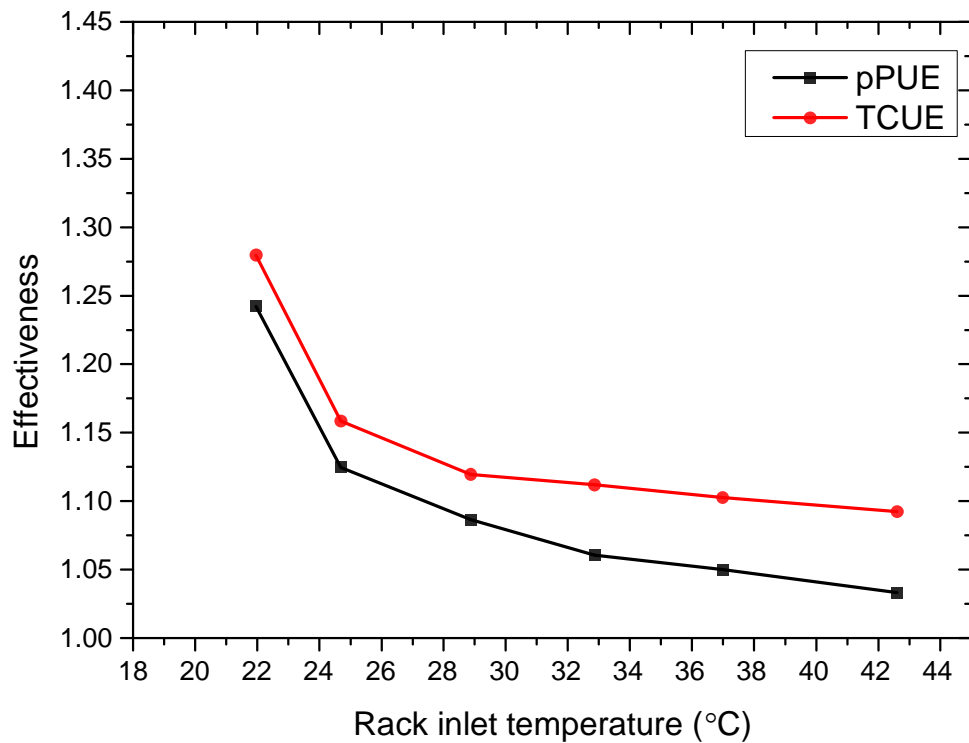
Figure 7.9 Effectiveness of the rack level as a function of the rack inlet water temperature for the 50% load operation.

7.5 High IT load of the rack

The high IT operational load includes the 75% and the 100% utilisation computational loads. The CPUs temperature is profoundly affected by the rack inlet temperature as shown in figure 7.10 and 7.11 for the 75% and 100% utilisation respectively. The average CPUs temperature increases from 44 to 66℃ when the rack inlet temperature increases from 22℃ to 44℃, respectively for the 75% utilisation load. The CPUs of two servers reached the critical temperature and shut down when the rack inlet temperature was 40℃. Under 100% utilisation, the CPUs temperature increased from 48 to 65℃ when the rack inlet temperature increased from 24 to 42℃ respectively. Here, four servers encountered shut down when the rack inlet temperature is higher than 42℃. Two of the servers' CPUs reached the critical temperature and shut down when the rack inlet temperature was 37℃.
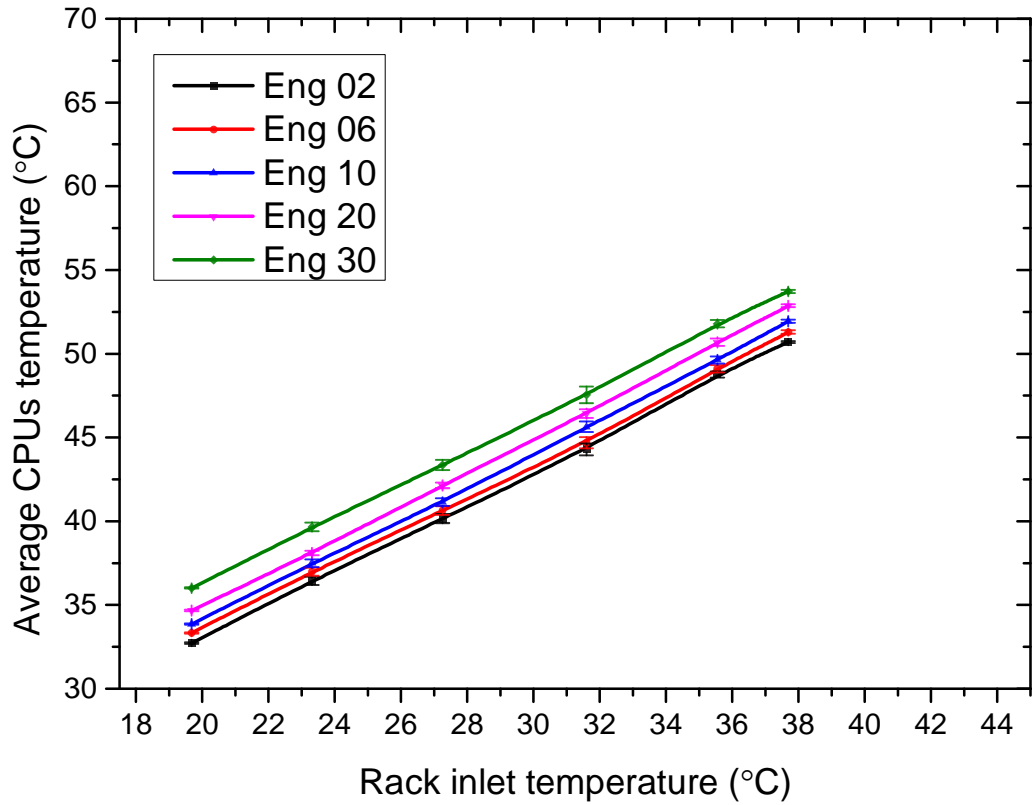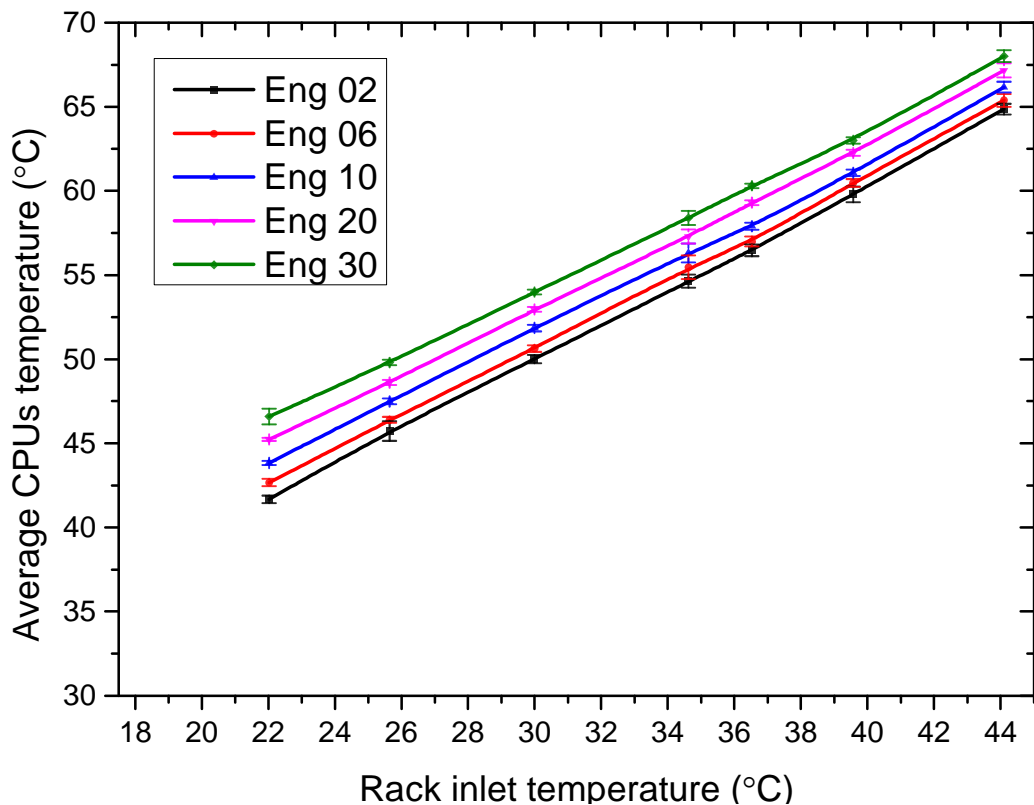

Figure 7.10 Average CPUs temperature of the five selected servers in the rack as a function of the rack inlet water temperature at 75% load
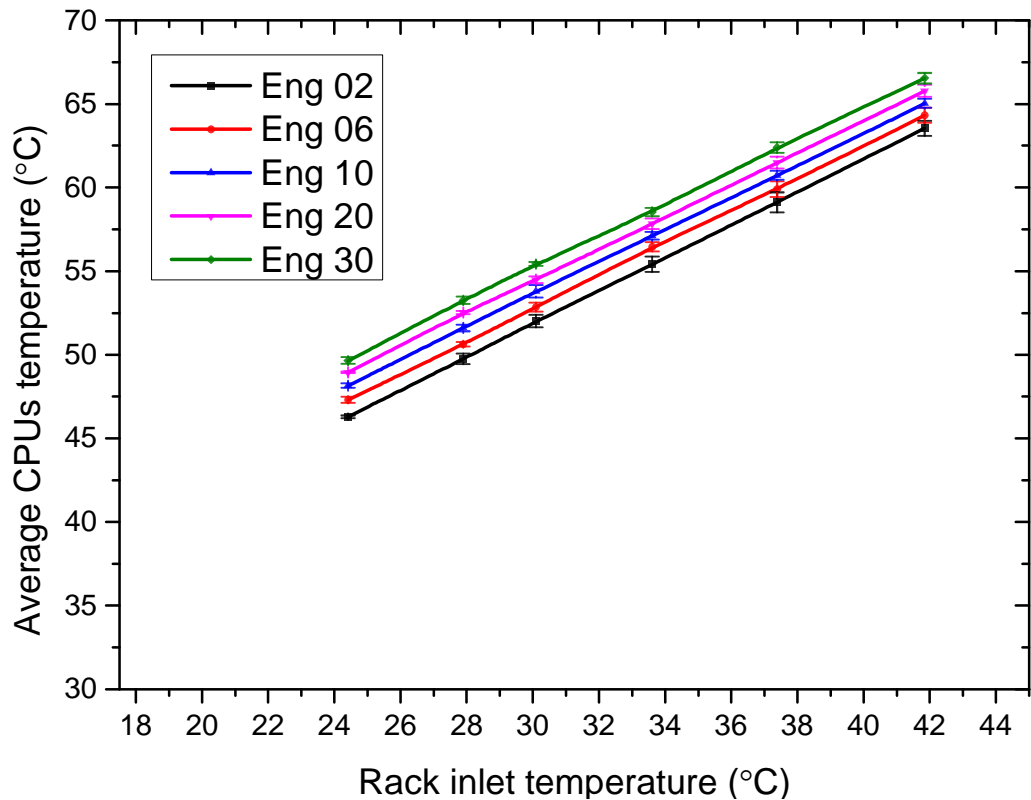
Figure 7.11 Average CPUs temperature of the five selected servers in the rack as a function of the rack inlet water temperature at 100% load

The average workload power efficiency was measured for the 75% and 100% utilisation experiments over various rack inlet water temperatures as shown in figures 7.12 and 7.13, respectively. The rack power consumption is also measured for the two stress levels. Figure 7.12 shows that the WPE of the rack reduced by 4% when increasing the rack inlet temperature from 22 to 44°C, while the power consumption of the rack increased by about 5% for this increase in temperature. On the other hand, for the 100% utilisation load, the decrease in the efficiency of the rack is reduced by 3% and the rack power consumption is increased by 4% when increasing the rack inlet temperature from 24 to 42°C
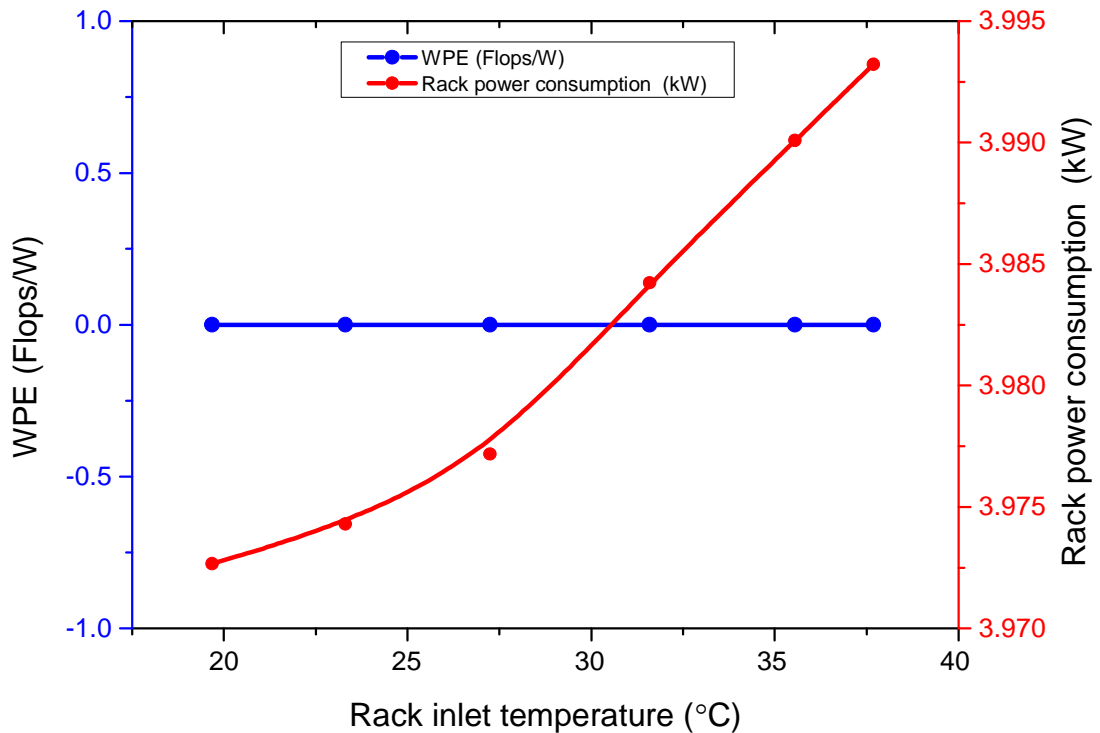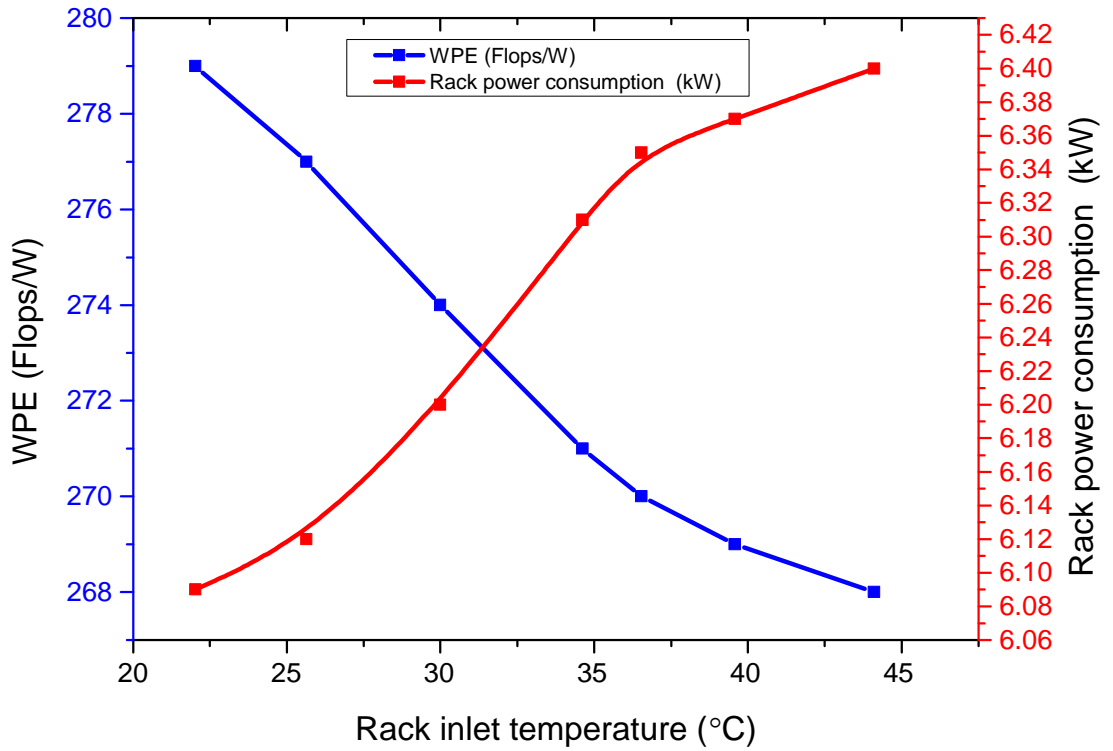
Figure 7.12 workload power efficiency and rack power consumption as a function of the rack inlet temperature at 75% load.
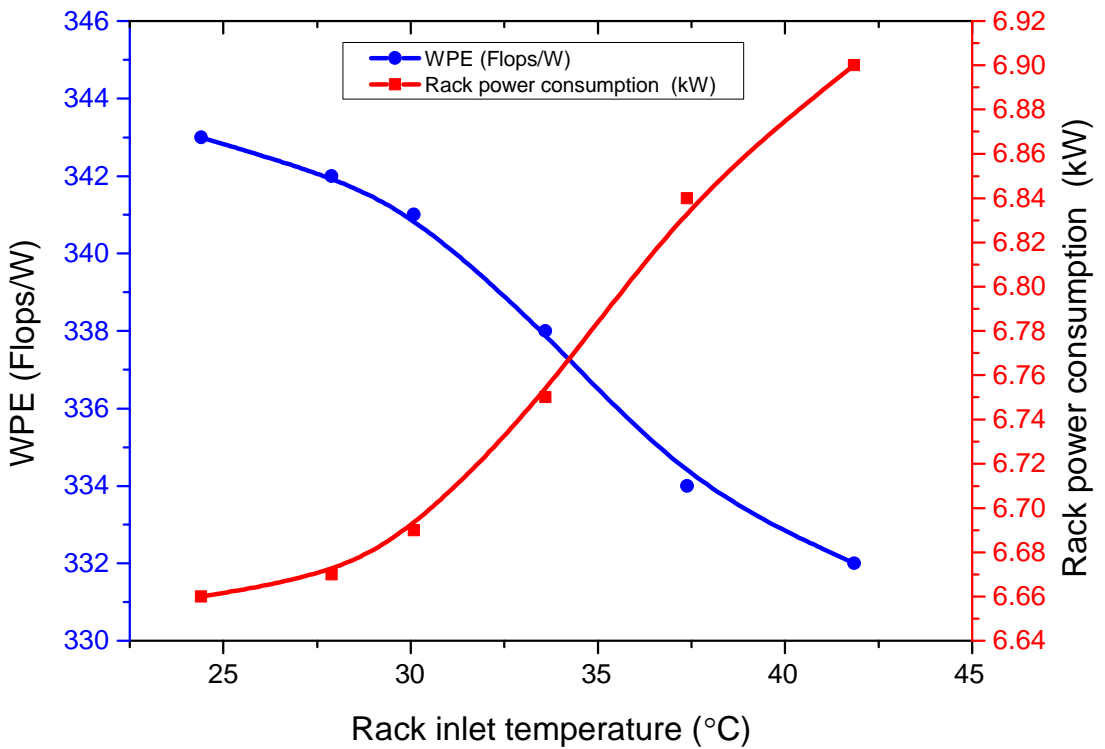
Figure 7.13 workload power efficiency and rack power consumption as a function of the rack inlet temperature at 100% load.

The effectiveness of the rack level is also measured in terms of the pPUE and the TCUE for the two stress level experiments of 75% and 100% utilisation as shown in figures 7.14 and 7.15, respectively. The pPUE is also improving with increasing rack inlet temperature as the power consumed by the AHU is decreasing and the power consumed by the rack is increasing for the two stress levels. The TCUE is also improving with increasing rack inlet temperature. However, the rate of improvement in the TCUE decreases at a rack inlet temperature of 30℃. The reason for this is that the fan speeds ramp up when the water inlet temperature to the rack is high and as the load level is high which means the CPUs do more work which makes it more sensitive to the inlet water temperature.

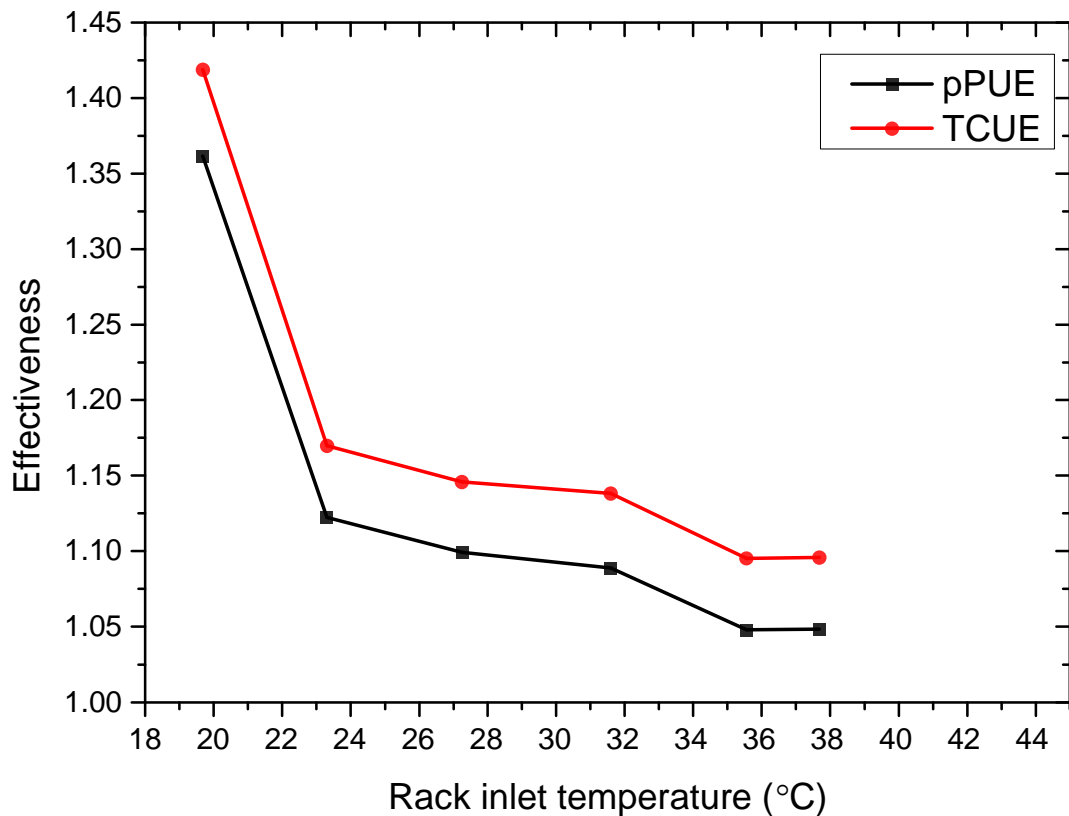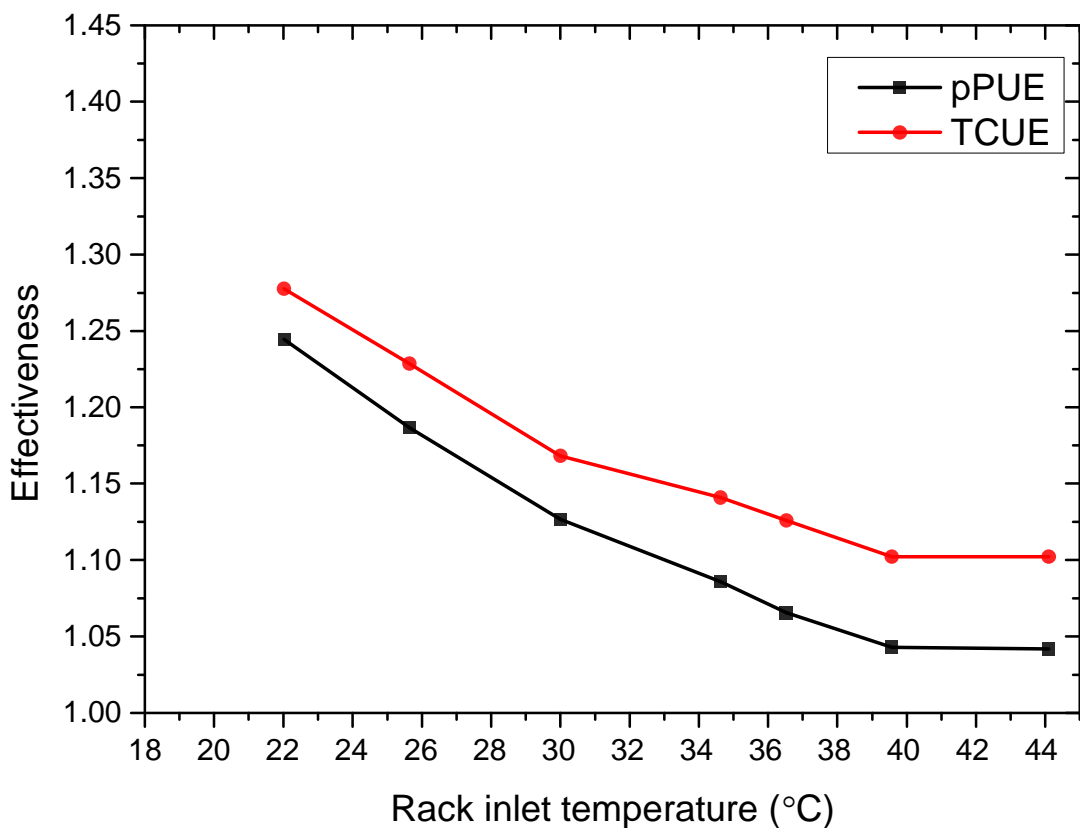**Figure 7.14 Effectiveness of the rack level as a function of the rack inlet water temperature for the 75% load operation.**
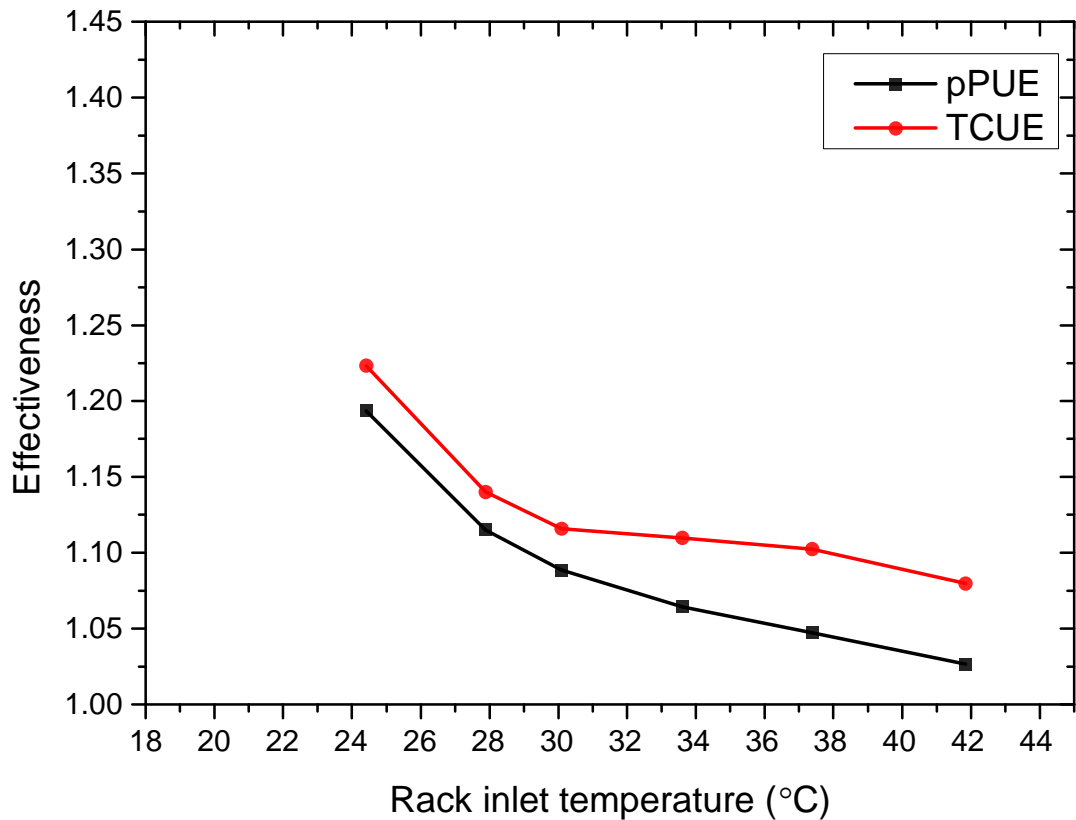
Figure 7.15 Effectiveness of the rack level as a function of the rack inlet water temperature for the 100% load operation.

7.6 Data centre workload power efficiency

The previous sections of this chapter have shown that the effectiveness of the data centre is improved when the rack inlet temperature increases as the work that needs to be done on the cooling side is reduced. This result is shown for the pPUE metric which is the most common metric used in the data centre industry. However, the improvement in data centre effectiveness (pPUE) is accompanied with a reduction in the computational performance of the servers as shown in the efficiency metric (WPE) and also with an increase in the rack power consumption. In addition, increasing the rack inlet temperature reduces the operational reliability and increases the problems in some of the servers modules and also reduces the robustness of the data centre [150].

Based on the results in the previous sections, it is necessary to choose the rack inlet temperature critically based on the utilisation load of the servers, and the ambient conditions. For this, and as mentioned in chapter 4, the

pDWPE metric which is bridging the gap between the effectiveness (pPUE) and the efficiency (WPE) is calculated in the present study to understand the relationship between these performance metrics and to help in choosing the best operating temperature to the rack. The pDWPE for the range of utilisation loads of the servers for various rack inlet water temperatures are shown in figures 7.16 and 7.17, respectively.

It can be seen from figure 7.16 that for the idle operation of the servers there is no pDWPE as the WPE is zero. It was also revealed in section 7.3 that the power consumption of the rack was altered slightly with increasing the rack inlet temperature. From this, it can be concluded that the rack inlet temperature can be as high as 38℃ with a significant improvement in the total energy consumption of the data centre and higher effectiveness. On the other hand, for the 25% and the 50% utilisation loads, the pDWPE is increasing with increasing rack inlet temperature. However, there is little change for inlet temperatures above 30℃.

The pDWPE for the experiments at high stress levels are shown in figure 7.17 for the 75% and 100% utilisation loads. As before, the pDWPE also improves with increasing rack inlet temperature. However, there is a high risk of equipment shut down when working with rack inlet temperatures higher than 35℃ as some of the server CPUs reach the critical temperature. From this analysis, working with operating temperatures of 30℃ for the high level stresses is ideal for high energy effectiveness, high efficiency and less total power consumption of the whole data centre.

From the above discussion, it can be concluded that the operating temperature of 30℃ can be the safest temperature to overcome some of the CPUs overheating as well as for better energy and operational performance of the rack of Sun Fire V20z servers of the current configuration.
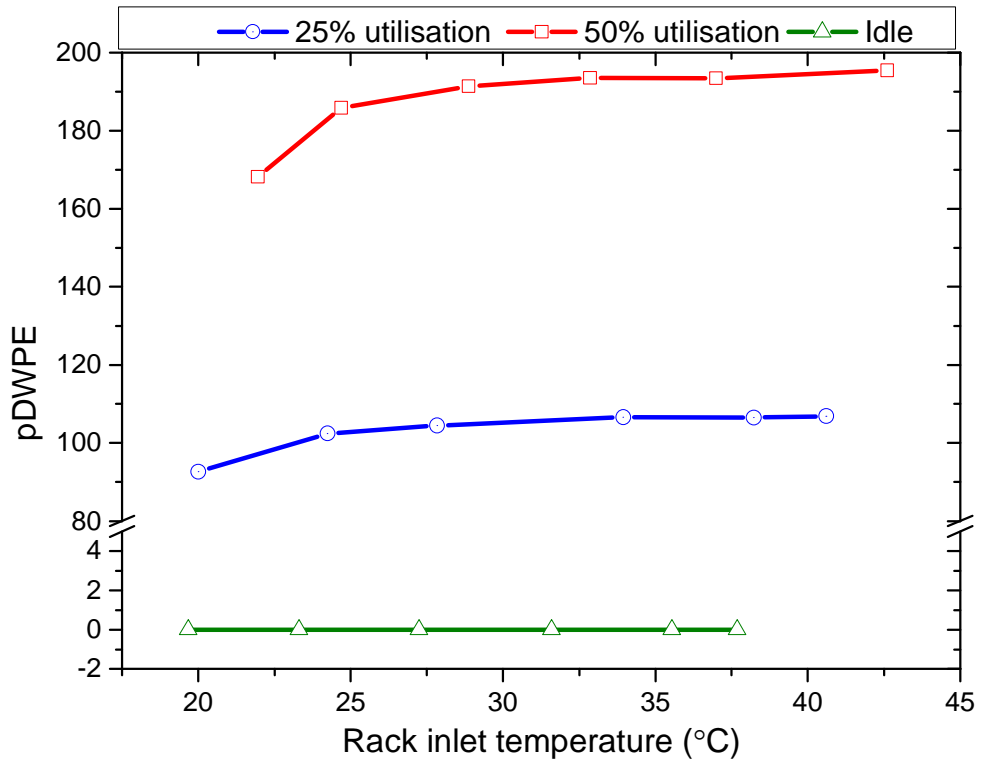
Figure 7.16 pDWPE as a function of the rack inlet water temperature for the idle and moderate loads.
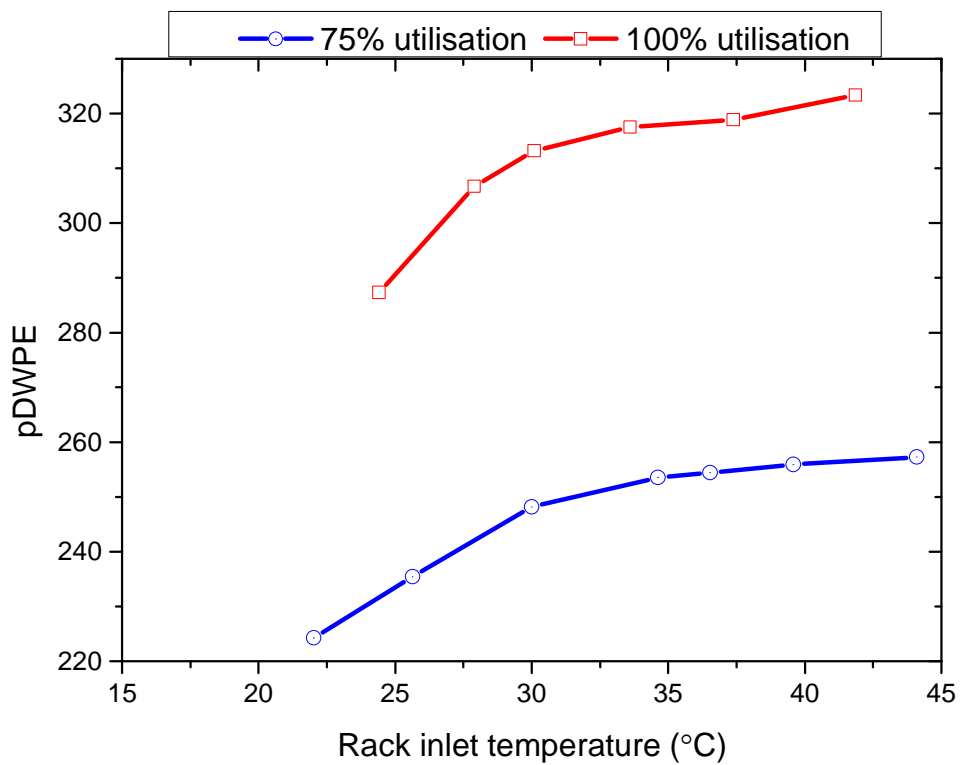
Figure 7.17 pDWPE as a function of the rack inlet water temperature for the high loads.

7.7 Summary

This chapter investigated the influence of rack inlet temperature on the CPU temperatures, the power consumption of the rack, the computing efficiency of the servers, the power effectiveness metrics and the compromise relation between the effectiveness and computational efficiency (pDWPE).

The CPUs temperature is strongly linked to the inlet coolant temperature to the rack. This influence is strongly affected by the computational load level. Some servers encountered shutdown when the inlet coolant temperature to the rack was high. At 50% utilisation load, two servers shutdown. The first shutdown happened at the inlet coolant temperature of 37℃ and the second shutdown at a temperature of 41℃. However, at the 100% utilisation load, four servers shutdown when the coolant inlet temperature increased to 37℃.

Increasing the inlet coolant temperature to the rack results in a high reduction in the power consumption of the AHU. On the other, the power consumption of the rack increases due to the higher current losses and higher fan speeds inside the servers with higher CPUs temperature. However, decreasing the power consumption of the AHU and increasing the power consumption of the IT have improved the effectiveness metrics (pPUE and TCUE) with increasing rack inlet temperature.

The computational efficiency of the servers was found to be adversely affected by increasing coolant inlet temperature. At 75% utilisation load, increasing the temperature inlet from 22 to 44℃ resulted in a 4% reduction in the WPE of the rack. Whilst at 100% utilisation, the WPE reduced by 3% when the rack inlet temperature was increased from 24 to 42℃.

Supplying high temperatures to the IT should be analysed critically as its benefits of energy savings of supplying less work on cooling should be weighed up against the draw backs of increasing rack power consumption and decreasing the computational efficiency as well as the detrimental effects on the operation of the servers such as server shutdown. Therefore, supplying coolant temperature of 30℃ is ideal for high IT load of the current servers and configurations.

Chapter 8 Droplet evaporation over heated surfaces

8.1 Introduction

The important role of the droplet evaporation phenomena in many applications has motivated many experimental and theoretical investigations. The importance of these applications is particularly well demonstrated in cooling systems [152-155] which form the link with the work presented in earlier chapters. Due to the high heat transfer associated with the phase change and evaporation process, work has been developed to show value in the use of evaporation and boiling in data centre cooling in general and in particular on chip cooling [2, 16]. Studies concerning droplet evaporation over heated surfaces explored the effect of the different parameters on the evaporation of droplets with particular focus on determining the effect of substrate thermal conductivity, roughness, and wettability on the evaporation process. A further link is with evaporation of liquid on heat exchangers surface where this exploits the high latent heat of evaporation. Single droplet studies are idealised since experience using spray cooling (chapter five) shows a deluge of water and film evaporation studies may be more pertinent [156-159]. Nevertheless, droplet studies do allow some of the underlying physics to be explored.

Droplet evaporation is still a challenging subject of discussions due to the associated complexity of fluid flow dynamics, the substrate physical surface chemistry and heat and mass transfer issues. Studies have explored the influence of substrate thermal conductivity on the evaporation lifetime of droplets [160-162], This has a key influence on the evaporation rate, for example David et al.[160] and Sobac and Brutin [162] have shown that there is a significant difference between the evaporation rate of nonmetallic and metallic substrates with the same hydrophilic properties while there are no major differences between different metallic substrates with similar hydrophilic properties. Surface wettability also has a key influence on the droplet evaporation processes. Both experimental and numerical studies have been developed to explain the interaction between droplet evaporation and wetting

characteristics of surfaces. The majority of these studies have concentrated on studying hydrophobic and super-hydrophobic surfaces [154, 162-166] while comparatively few have investigated droplet evaporation over hydrophilic substrates [162, 163]. Studies concerning free droplet evaporation (no substrate temperature control) have received more attention [154, 163-165, 167-171] than evaporation under controlled substrate temperatures [162, 166, 167].

Characterising surface properties and the interactions between the droplet and the substrate is necessary for numerical investigations. Deegan et al. [172] proposed a model to predict the local evaporation flux over the droplet surface as a function of the contact angle. Popov [173] developed a model to calculate the total rate of evaporation based on the local evaporation flux model of Deegan et al. which is called the vapour-diffusion model. This model underpins many studies for different surface wettability spanning hydrophilic [162, 172], pinned hydrophobic [162], de-pinned hydrophobic [166, 174] and super-hydrophobic surfaces [164]. The model predicts the evaporation rate with a reasonable accuracy for the hydrophobic [166] and hydrophilic [162, 166] surfaces while the model was found to significantly under predict the total time of evaporation for the super-hydrophobic surfaces [166, 174]. Dash and Garimella [174] proposed a correction factor for the vapour-diffusion model to reduce the discrepancy of the predicted results for the super-hydrophobic surfaces. Gleason and Putnam [175] introduced a modification for the local mass flux and the evaporation rate models by applying a temperature distribution across the liquid-vapour interface of the droplet. Pan et al. [176] attributed the discrepancy of the vapour-diffusion model for the evaporative cooling while Carle et al. [177] claimed this underestimation to be due to the natural convection driven by the buoyancy in the gas phase for the heated surfaces. However, others have attributed the underestimation to the effect of substrate thermal conductivity which is not included in the model [162, 164, 178].

Two main evaporation modes have been discovered by experimental observations which are related to the surface properties and characterised by

the pinning and de-pinning properties of the droplet contact line: (i) constant contact radius mode (CCR), in which the contact angle of the droplet vanishes; (ii) the constant contact angle mode (CCA), in which the contact radius of the droplet vanishes [179]. However, another phase of evaporation have also been reported to occur shortly in some cases at which both of the contact angle and the contact radius collapse [171, 180]. The forces associated with the role of evaporation mode depends mainly on the surface roughness[181], wettability [182] and surface energy [183]. Mollaret et al. [161] found that for high surface energy the pinning force increase with increased substrate temperature while there is a little dependence on temperature for the low energy surfaces. Blake and Coninck [184] and Putnam et al. [182] both examined the effect of surface wettability on the pinning and de-pinning forces and found that the pining force is more dominant for the contact angles of lower than 90° (hydrophilic surfaces) and reduces with increasing the contact angle (hydrophobic surfaces). Therefore, hydrophobic surfaces are associated with the CCA evaporation mode [174, 185], although wettability is not the only defining factor for all the hydrophobic surfaces [164]. Kulinich and Farzaneh [186] determined that the mode of evaporation can be discovered from the contact angle hysteresis of the surface and cannot be explained by the initial contact angle of the droplet.

The mode of evaporation is vital in determining the evaporation mass flux and the total time of evaporation. Models focus either on a single mode of evaporation (CCR or CCA) that has the largest portion of the total evaporation time based on the experimental visual observation [162, 166, 174], whilst others have considered using both of the models (CCR and CCA) to describe the evaporation process where the droplet evaporates at the CCR mode and then shifts to CCA mode based on the experimental data where both of the contact angle and the droplet radius are visually measured and monitored to be used to identify the point at which the modes of evaporations are shifted [154, 171, 187, 188]. Some studies have developed an experimental control methodology (such as creating a trench in the substrate) to force the droplet to evaporate at constant contact radius allowing theoretical predictions for a

specific model of operation [175]. However, these aforementioned models are essentially based on the experimental visual observations of all the droplet evaporation processes on the studied surfaces to be depicted and applied in the proposed evaporation model. Stauber et al. [187] proposed a relationship between the initial contact angle and the point at which the droplet shifts into the de-pining mode based on the experimental data of droplet evaporation from different studies in the literature. Recently, Hu and Wu [188] studied theoretically the effect of different receding contact angles on the total time of evaporation assuming the droplet evaporates under pinned mode and then in the de-pinning mode once the receding contact angle was reached.

The majority of studies have considered evaporation of droplets with volumes between $1\mu l$ and $5\mu l$ [154, 162, 164-166, 170], since these are consistent with the assumption of a spherical droplet cap where surface tension dominates over gravity to yield a homogeneous and regular droplet shape. This then allows a visualization method to be adopted based on imaging the droplet from the side in order to determine the contact angle, contact radius, and droplet height [154, 160-171, 174, 189, 190]. The Young-Laplace approach, valid for axisymmetric droplets, is the most common way to fit the data. However, this approach has limitations since estimates of the contact line position is a very sensitive process which is associated with large errors in the predicted contact line and contact angle [191]. The associated error increases at small contact angles and rates of droplet evaporation are often extrapolated from data where de-pinning occurs [160, 162, 178, 192]. There is an additional error at large droplet sizes associated with a non-circular contact line [164]. Moreover, Srinivasan et al [191] have shown that a small variation in the lighting and focus conditions can also lead to a difference of several degrees in the measured contact angle which in turn leads to significant errors in the calculated droplet volumes.

In this chapter, the evaporation of sessile droplets over heated surfaces is examined experimentally and theoretically using the vapour-diffusion model. The originality of the present work is to:

1. Propose a new way of measuring the droplet evaporation using a precise measurement balance to overcome the limitation through the conventional imaging methods. A self-contained compact and light weight (50 grams) experimental heating stage is designed with a PID controller to control the substrate surface temperature. The light weight component allows the use of a 10-microgram resolution balance with a maximum stage load of 220 grams to study droplet evaporation over a large droplet size range (2 μl to 30 μl) for hydrophobic and hydrophilic surfaces.

2. Postulate a new method of using the theoretical model to predict the evaporation rates over the whole period of drying with a switch from CCR to CCA modes (stick-slip mode) based on a measured receding contact angle. As previously discussed, the use of the stick-slip mode of evaporation reported in the literature has been limited to experimental observation of the whole evaporation process.

8.2 Experimental setup

The principle of the experiment consists of deposition of a droplet of a controlled volume on a heated substrate, which is then allowed to evaporate. The time evolution of the weight of the droplet is monitored and measured.

To investigate the influence of the sample substrate temperature, a compact light weight (50 grams) cylindrical container was designed to carry the substrate sample with the electronic circuit to heat and control the temperature as shown in figure 8.1. The substrate is heated using a thick film electrical resistor powered by a small LiPo battery and controlled using an embedded PID controller. A microcontroller (Teensy 3.1) was used to build a PID algorithm (Appendix A) which is designed and programmed to control the power input to the thermal resistor based on the measured temperature in the substrate as shown in figure 8.2. A K-type thermocouple is connected to an analogue output amplifier (Adafruit AD8495) which in turn is connected to the microcontroller. The temperature sensor was calibrated against a commercial

temperature measurement system in a range of temperatures from 0 to 100℃. The microcontroller is connected to a Mosfet (FQP30N06L) which in turn allows the current to flow to the thermal resistor. The amount of current passing across the Mosfet to the thermal resistor is proportional to the incoming voltage signal from the microcontroller. This signal is an analogue signal which is converted from a digital signal in the microcontroller using Pulse Width Modulation (PWM). All of the electronic components are powered by a compact light weight polymer lithium ion battery (3.7V, 1Ah). The circuit enables the substrate temperature to be kept constant with a variation of ± 0.5℃ by modulating the amount of current to the resistor. Using the PID controller loop means that any temperature disturbance (e.g. reduced heat load during drying) will be automatically compensated for. The substrate under test is attached to the thermal resistor with a thin uniform layer of thermally conductive ethoxy paste (Electrolube TCER75S). For every temperature set point, the code was modified and re-uploaded. After uploading the code, the temperature in the substrate sample was monitored until it reached a steady state at the set temperature point. To ensure for the reliability of the temperature controller, the temperature is also monitored by an external thermocouple via the hole that was designed to be extended to the outside of the container (as shown in figure 8.1 (point 4)). After that the circuit is packed into the designed container to be weighted in the balance.

An analytical balance (METTLER TOLEDO XP 205) is used to measure and digitally record the droplet mass evolution as a function of time with a resolution of 10 micrograms. The design of the heated stage allows the self-contained unit to sit on the weighing pan, with sides around the balance to prevent air currents. The balance is tared for every experiment and connected to a computer via the Real Term software and the weight is recorded every half a second.

A camera (Basler acA1300-30μm and lens CCS SE-16SM) is used to obtain side and top views of the droplet. The recorded frames allow monitoring of the droplet lifetime as well as the shape of the droplet especially at the last stage of the droplet life.

The droplet is created using a micropipette to control the droplet size, which is gently laid down on the plate sample to evaporate into the air inside a cubic cell of dimensions $(15 \times 18 \times 21\text{cm}^3)$ to avoid any perturbation from potential external flow. To ensure repeatability of the results, each experiment is repeated four times. Before each measurement, a cleaning procedure is adopted by rinsing the substrate with deionized water and drying by a blower.

The temperature and relative humidity inside the room containing the balance are monitored using local room control and maintained at $21\pm0.2°\text{C}$ and $30\pm1\%$, respectively. The effect of substrate temperature was investigated over five different temperatures: 40, 45, 50, 55 and 60°C (with an error of $\pm0.5°\text{C}$).

Deionized water droplets of four different initial volumes (2, 4, 8 and $30\mu l$) are investigated in the experiment. The range of droplet volumes over droplet shapes dominated by surface tension through to these dominated by gravity; the Bond number (B_o) [193] varies from $B_o < 1$ and $B_o > 1$, respectively. The Bond number is defined by:

$$B_o = \frac{\Delta \rho g R^2}{\gamma} \qquad\qquad 8.1$$

Where $\Delta\rho$ is the density difference between the water and air, g is the gravitational acceleration, γ is the surface tension between water and air at the substrate temperature, and R is the droplet contact radius.
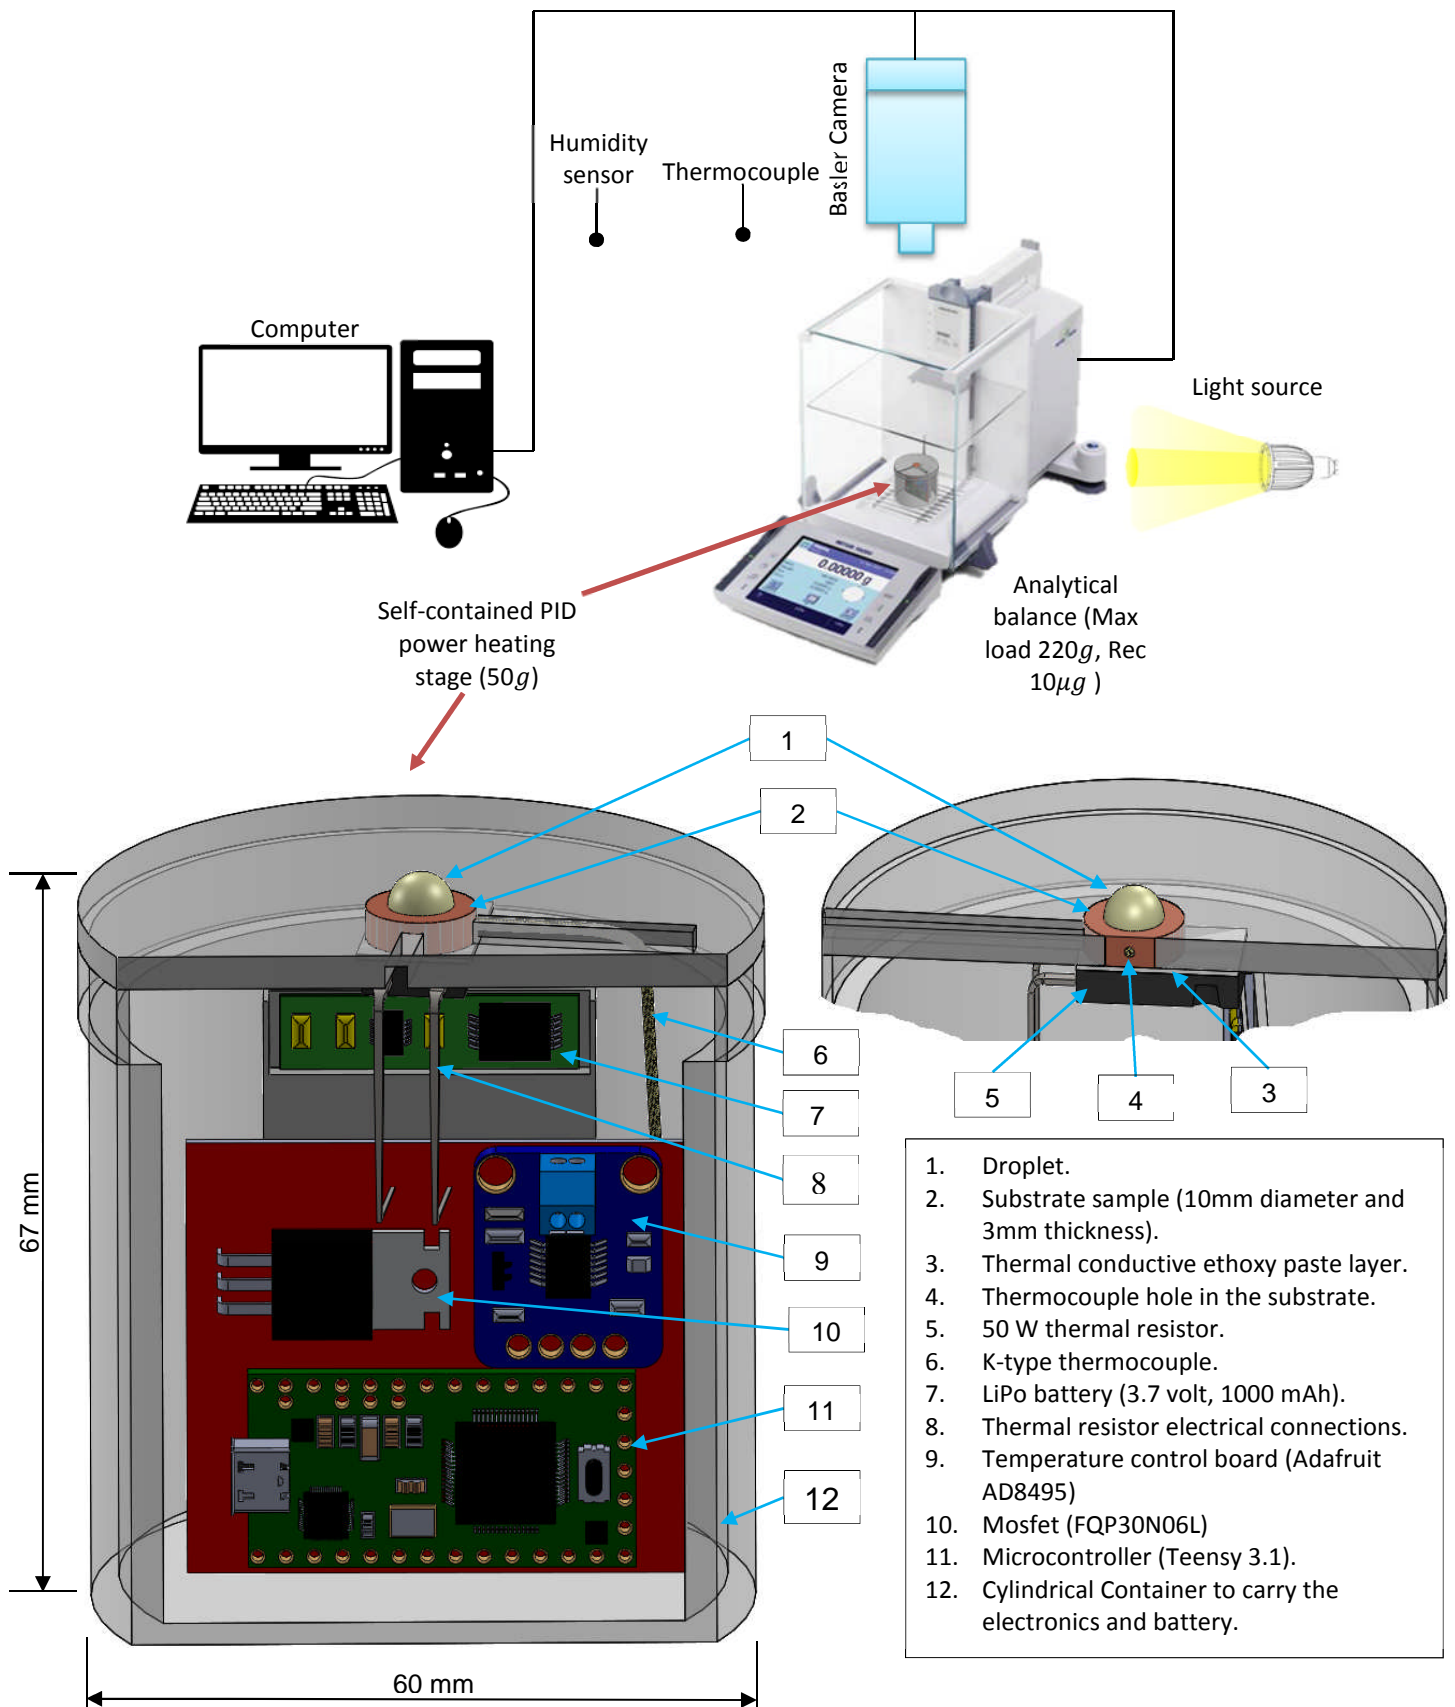
Figure 8.1 Experimental rig setup to track droplet evaporation.
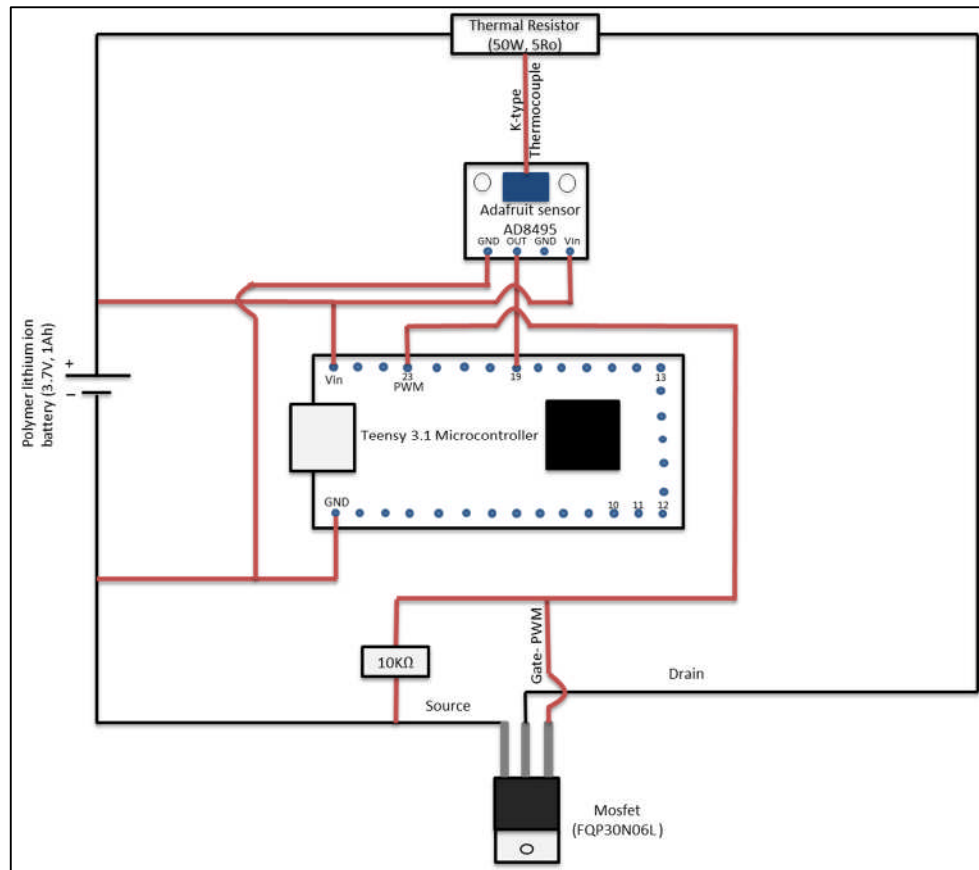
Figure 8.2 PID electronic circuit temperature controller.

8.3 Sample preparation and characterization

Two different substrates were studied, hydrophobic treated aluminium and hydrophilic treated aluminium. Disks of material were formed by machining an aluminium bar of 10mm diameter into disks of 3mm thickness. The hydrophobic coated disks were polished from both sides with an emery paper of 1200 grade before being coated using Oxford nanoSystems [194] nanoFLUX® process to produce a dendritic surface structure of nano sizes and morphologies. Each disk has two holes (from the sides) for temperature measurement, one of these holes is used to embed the electronic circuit thermocouple sensor whereas the other one is used for temperature calibration purposes (see figure 8.1, label 4).

226

8.3.1 Contact angle and receding contact angle measurements

The contact angle for all the samples was measured using CAM 200 from KSV. Each measurement was taken as the average of five different measurements after dispensing the droplet to ensure the dynamic forces of dispensing do not affect the reading. This was repeated at four locations across the surface. The static contact angles of the hydrophobic and hydrophilic aluminium surfaces were 120°±1.83 and 49°±1.13, respectively.

To measure the receding contact angle, the First Ten Angstroms (FTA) 4000 Microdrop is used. This method is a video droplet shape analysis system for small droplets from picolitres up to microlitre regimes. It automates the vertical motion to control the dispense tip and the droplet volume. The optimal height of the tip can be adjusted using real time image processing to be positioned above the sample surface to touch off the droplet. Two microscopes are fitted to record the frames, which are the primary and secondary microscopes as shown in figure 8.3. The primary microscope is an analytical microscope which has a zoomed horizontal field of view from 2mm to 200 microns to capture the side views of the droplet which is back lit by a high intensity LED source. Whereas the secondary microscope is to view the droplet and the substrate sample from 60° angle, which is provided with incandescent light from the top side [195].

The measurement procedure is first started by growing water droplets onto the sample using the provided piezo pump and letting it stabilize for 20s to release the droplet from stresses. To ensure the reliability of the results, three different droplet sizes are chosen. Then, the needle is inserted inside the drop either by changing the vertical position of the needle or by changing the vertical position of the sample. After that, the droplet size is then reduced slowly by sucking the droplet liquid using the piezo pump until the droplet contact radius just starts moving. The video recording commencing at an appropriate timing where many frame captures should be taken covering the point at which the droplet contact area starts moving, representing the receding contact angle. Finally, the Non-Spherical fitting method is used with

manual baseline determination to acquire the receding contact angle values which are calculated as the average between the left and right angles.
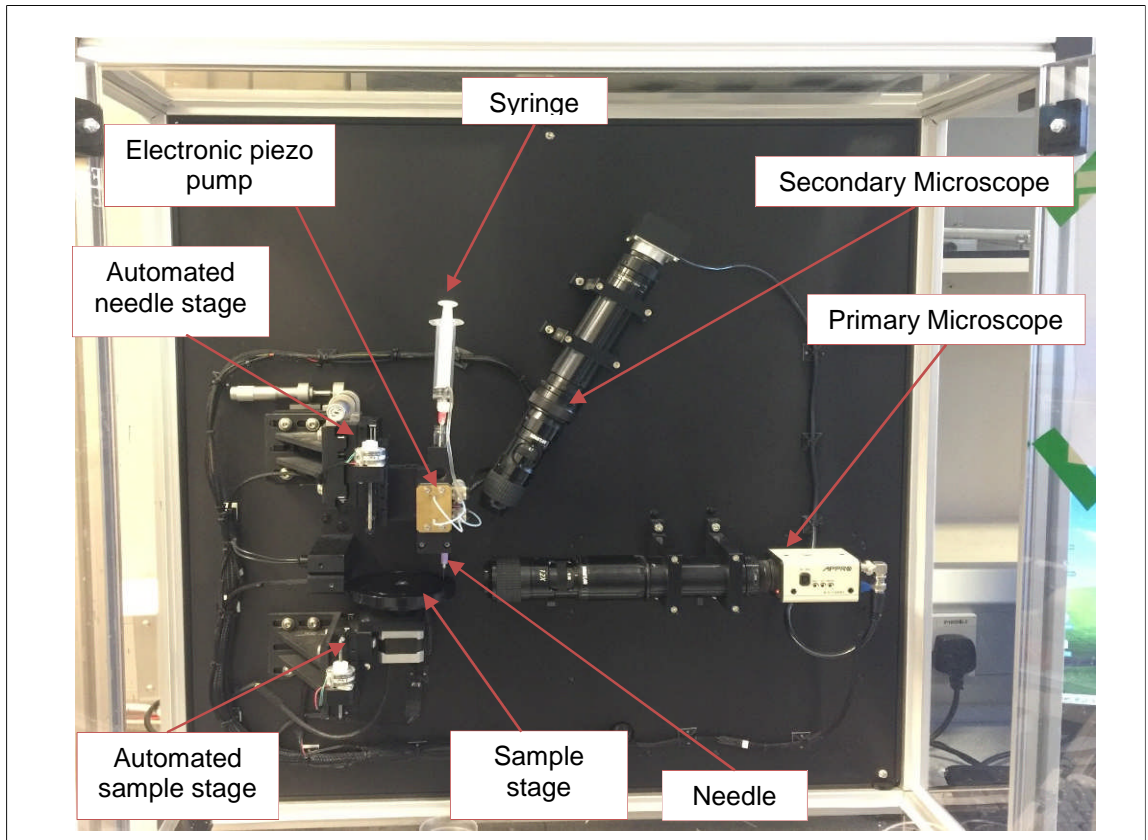
Figure 8.3 FTA4000 microdrop rig for measuring the receding contact angle.

The measured receding contact angle was found to be $10° \pm 0.6$ for the aluminium hydrophobic surface. For the aluminium hydrophilic surface, the contact line starts receding at a value of $6° \pm 2$. The high error is due to the strong pinning force in the hydrophilic surfaces as was observed previously in the literature [182, 184]. Another error source is the precision of the instrumentation, where higher errors are associated with small receding contact angles found on hydrophilic surfaces.

8.3.2 Surface topography of the substrates

The surface roughness of the samples were determined using white light interferometry (Bruker Np-Flex). The 3-D images of the solid surfaces were obtained over a region of $0.838 \times 1.12\ mm^2$ as shown in figures 8.4, and 8.5. The root mean square roughness (RMS) was found to be 2.468 and 0.677μm

whereas the arithmetic mean height (Sa) values were 1.9 and 0.238 µm for the hydrophobic and hydrophilic aluminium surfaces, respectively.
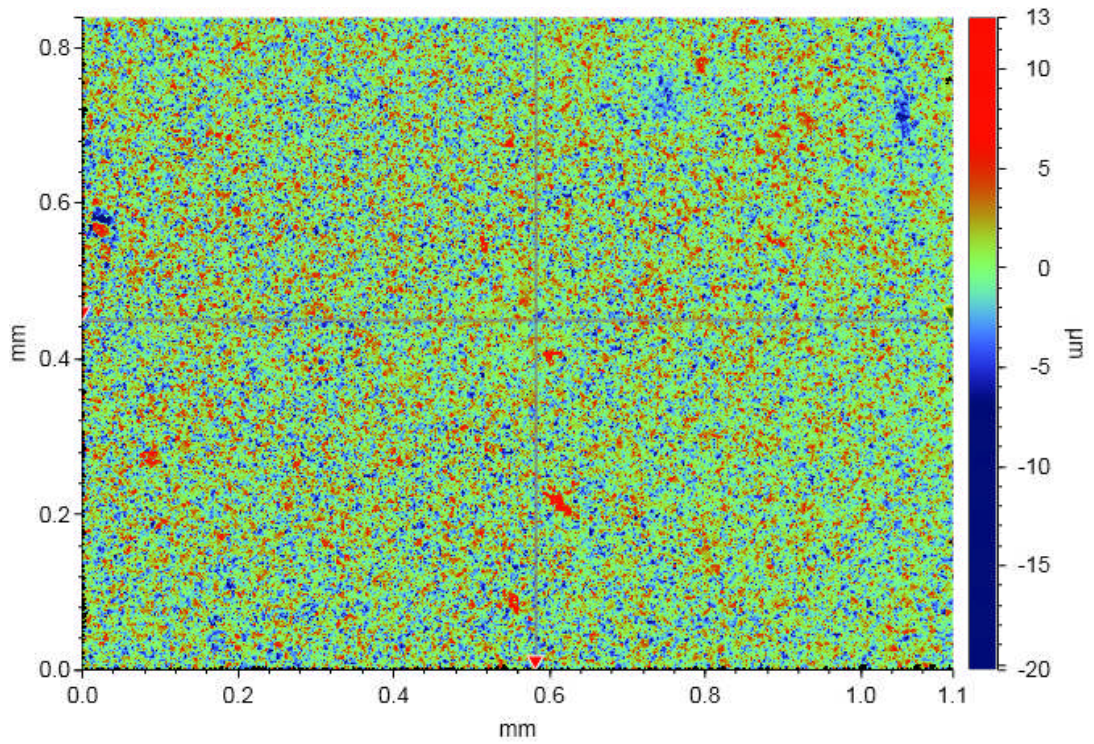
Figure 8.4 Aluminium-hydrophobic coating surface.
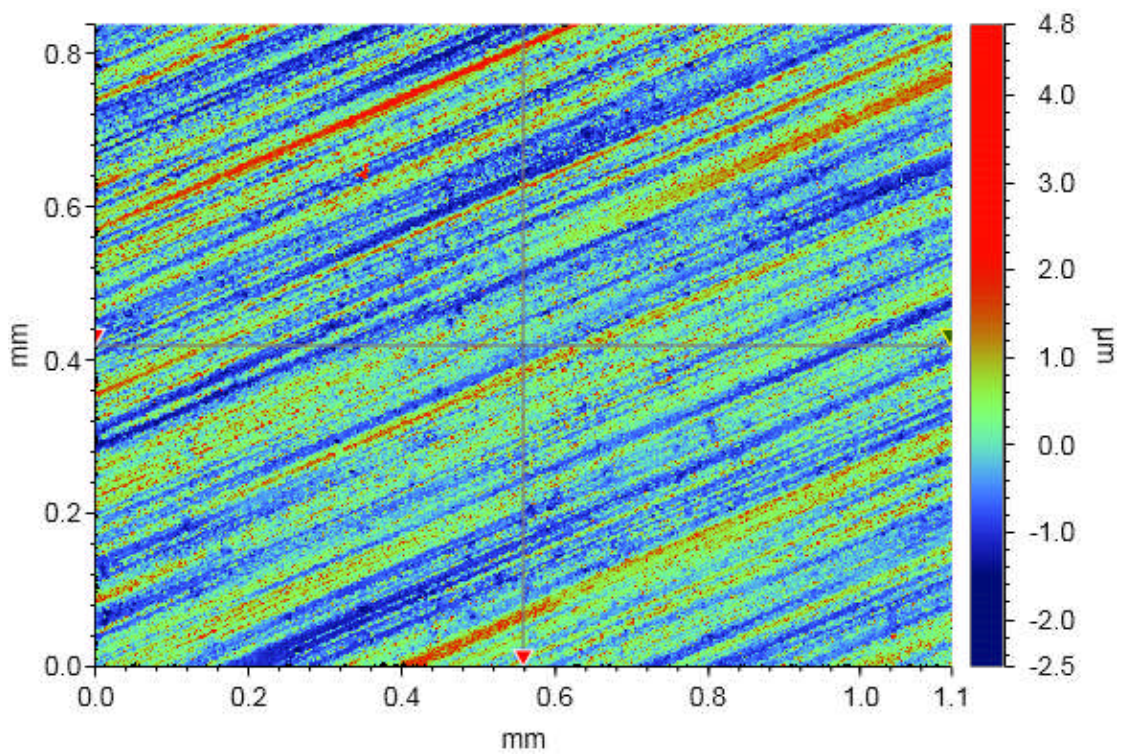
Figure 8.5 Aluminium-hydrophilic coated surface.

8.4 Theoretical analysis

The theoretical analysis presented here predicts the evaporation of a sessile droplet due to forced evaporation through heating of the substrate, allowing comparison with experimental data.

In the diffusion driven evaporation model, the limiting rate of mass transfer between the liquid and vapour medias is due to diffusion at the liquid-vapour interface. The second Fick's law of diffusion governs the evaporation of the microdroplet

$$\frac{\partial c}{\partial t} = D\nabla^2 c \qquad\qquad 8.2$$

where, D, c and t are the diffusion constant, vapour concentration and time respectively.

The time scale analysis indicates that the droplet evaporation can be assumed to be as quasi-steady when the droplet lifetime is much larger than the mass transfer as indicated by Fourier number (D/R_i^2), where R_i is the initial droplet radius. In this case, the transient term in Fick's law can be neglected, resulting in the Laplace's equation.

$$\nabla^2 c = 0 \qquad\qquad 8.3$$
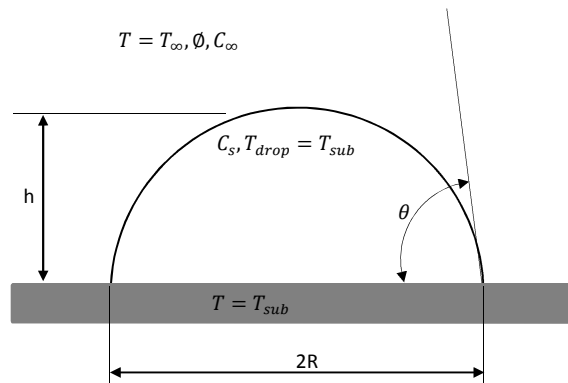
Figure 8.6 Droplet schematic and the boundary conditions.

The droplet size adapted in this theoretical analysis is considered to satisfy the spherical cap assumption and to indicate symmetry across the central axis of the droplet. The boundary conditions (as shown in figure 8.6) to satisfy Laplace's equation are (i) saturated vapour at the surface of the droplet at the substrate temperature (c_s), (ii) the vapour concentration far from the droplet surface is defined by the ambient conditions ($c_\infty = \emptyset c_{\infty,s}$), where $\emptyset$ is the relative humidity and $c_{\infty,s}$ is the saturated vapour concentration at the ambient temperature, (iii) all the phase change occurs across the liquid-vapour interface of the droplet, and (iv) the diffusion coefficient is calculated based on the substrate temperature. By transforming the domain into toroidal coordinate system as explained by Lebedev [196], Laplace's equation can be solved for the vapour concentration around the droplet and the solution can be described as

$$\frac{c(\propto-\beta)-c_\infty}{c_s-c_\infty} = \sqrt{2\cosh\propto -2\cos\beta} \times \int_0^\infty \frac{\cosh(\theta\tau)\cosh(2\pi-\beta)\tau}{\cosh(\pi\tau)\cosh(\pi-\beta)\tau} P_{\frac{1}{2}+i\tau}(\cosh\alpha)d\tau \qquad 8.4$$

Where $P_{\frac{1}{2}+i\tau}(\cosh\alpha)$ is the hyperbolic Legendre function.

From equation (8.3) and (8.4), the Deegan model can be obtained which governs the local evaporation flux [172, 173]

$$J(\propto) = \frac{D(c_s-c_\infty)}{R}\left[\frac{1}{2}\sin\theta + \sqrt{2}(\cosh\alpha+\cos\theta)^{3/2} \times \int_0^\infty \frac{\tau\cosh(\theta\tau)}{\cosh(\pi\tau)}\tanh[(\pi-\right.$$

$$\left.\theta)\tau]\, P_{i\tau-\frac{1}{2}}(\cosh\alpha)d\tau\right] \qquad 8.5$$

The total mass loss of the droplet during evaporation can be obtained by integrating the Deegan model, equation (8.5), over the liquid-vapour interface which results in Popov's model [173, 174]

$$\frac{dM}{dt} = \rho_L\frac{dV}{dt} = -\pi RD(c_s - c_\infty)f(\theta) \qquad 8.6$$

$$f(\theta) = \frac{\sin\theta}{1+\cos\theta} + 4\int_0^\infty \frac{1+\cosh 2\theta\tau}{\sin 2\pi\tau}\tanh[(\pi-\theta)\tau]\,d\tau \qquad 8.7$$

where M, ρ_L, V, and R are the mass, density, volume, and contact radius of the droplet respectively.

The rate of mass change of the droplet of spherical cap can be calculated as a function of the base radius and contact angle. Two stages of evaporation are observed experimentally in which the major part of evaporation occurs when the droplet base radius, R, is constant and the contact angle is a function of time, $\theta = \theta(t)$, until it reaches the receding contact angle, $\theta = \theta_r$, where the droplet starts de-pining and the evaporation becomes a function of base radius only, $R = R(t)$ with the contact angle remaining fixed at the receding angle until complete evaporation. The evolution of the contact angle during the first stage of evaporation of a pinned droplet and the second stage of evaporation of de-pinned droplet can be respectively represented as

$$\frac{d\theta}{dt} = \frac{D(c_s - c_\infty)}{\rho R_i^2}(1 + \cos\theta)^2 f(\theta) \qquad 8.8$$

$$\frac{dR}{dt} = \frac{D(c_s - c_\infty)}{\rho R^2}\frac{(1+\cos\theta_r)^2}{2+\cos\theta_r}\left[\frac{1}{1+\cos\theta_r} + 4\int_0^\infty \frac{1+\cosh 2\theta_r\tau}{\sin\theta_r \sinh(2\pi\tau)}\tanh[(\pi - \theta_r)\tau]\,d\tau\right] \qquad 8.9$$

The new mass of the droplet after each time step based on the new contact angle or the new base radius is calculated based on the spherical cap assumption as $(M = \rho_L V)$

$$M = \frac{\pi\rho R^3}{3}\frac{(1-\cos\theta)^2(2+\cos\theta)}{\sin^3\theta} \qquad 8.10$$

The transient contact angle in equation (8.8) and the transient base radius in equation (8.9) were numerically integrated using the Trapezoidal rule method using Matlab [197]. The calculations procedure is explained in figure 8.7. The diffusion only model calculations are solved using the MATLAB code. The code was written to enable the user to choose the simplified assumption on which Popov model can be solved. These assumptions can be based on the CCR mode or the CCA mode or the SSR mode. In any assumptions used, the initial and boundary conditions are needed to be fed into the code. These conditions are: initial droplet radius (R), ambient relative humidity $(\emptyset)$, ambient temperature (T_∞), substrate temperature (T_s), water density (ρ), initial contact angle (θ), and the receding contact angle (θ_r).

The latent heat of evaporation (h_{fg}) is linearly decreasing with temperature and can be calculated from the following relationship for the range of temperatures in the present experiments [198]:

$$h_{fg} = -2.4T + 2502.8 \qquad\qquad 8.11$$

Where T is the temperature in degree Celsius.

While the saturated vapour pressure can be calculated using the Clausius–Clapeyron equation [199].

$$ln\frac{P}{P_{ref}} = \frac{Mh_{fg}}{R_g}\left(\frac{1}{T_{ref}} - \frac{1}{T}\right) \qquad\qquad 8.12$$

Where: P is the saturation pressure, T is the temperature at the required saturated vapour pressure, P_{ref} is the known reference pressure at the reference temperature (T_{ref}), M is the molar mass of water vapour and R_g is the universal gas constant.

The ideal gas law is used to calculate the saturated vapour concentration (c) at the droplet surface and the ambient as follows:

$$c = \frac{MP}{RT} \qquad\qquad 8.13$$

The droplet surface concentration c_s is considered to be saturated and can be calculated based on the substrate temperature T_s. While the ambient vapour concentration c_∞ is calculated based on the ambient temperature T_∞ and the relative humidity of the surrounding air.

The diffusion coefficient of water vapour into air varies with the substrate temperature. In most of the studies, this value is assumed constant in the droplet modelling. However, in the present study where the range of temperatures are between 313K and 333K, the Marrero and Mason empirical correlation is used [200, 201]:

$$D = 1.87 \times 10^{-10}\frac{T_s^{2.072}}{P_{atm}} \qquad\qquad 8.14$$

Where (P_{atm}) is the ambient pressure in atmosphere (1atm), and the substrate temperature T_s is in kelvin.
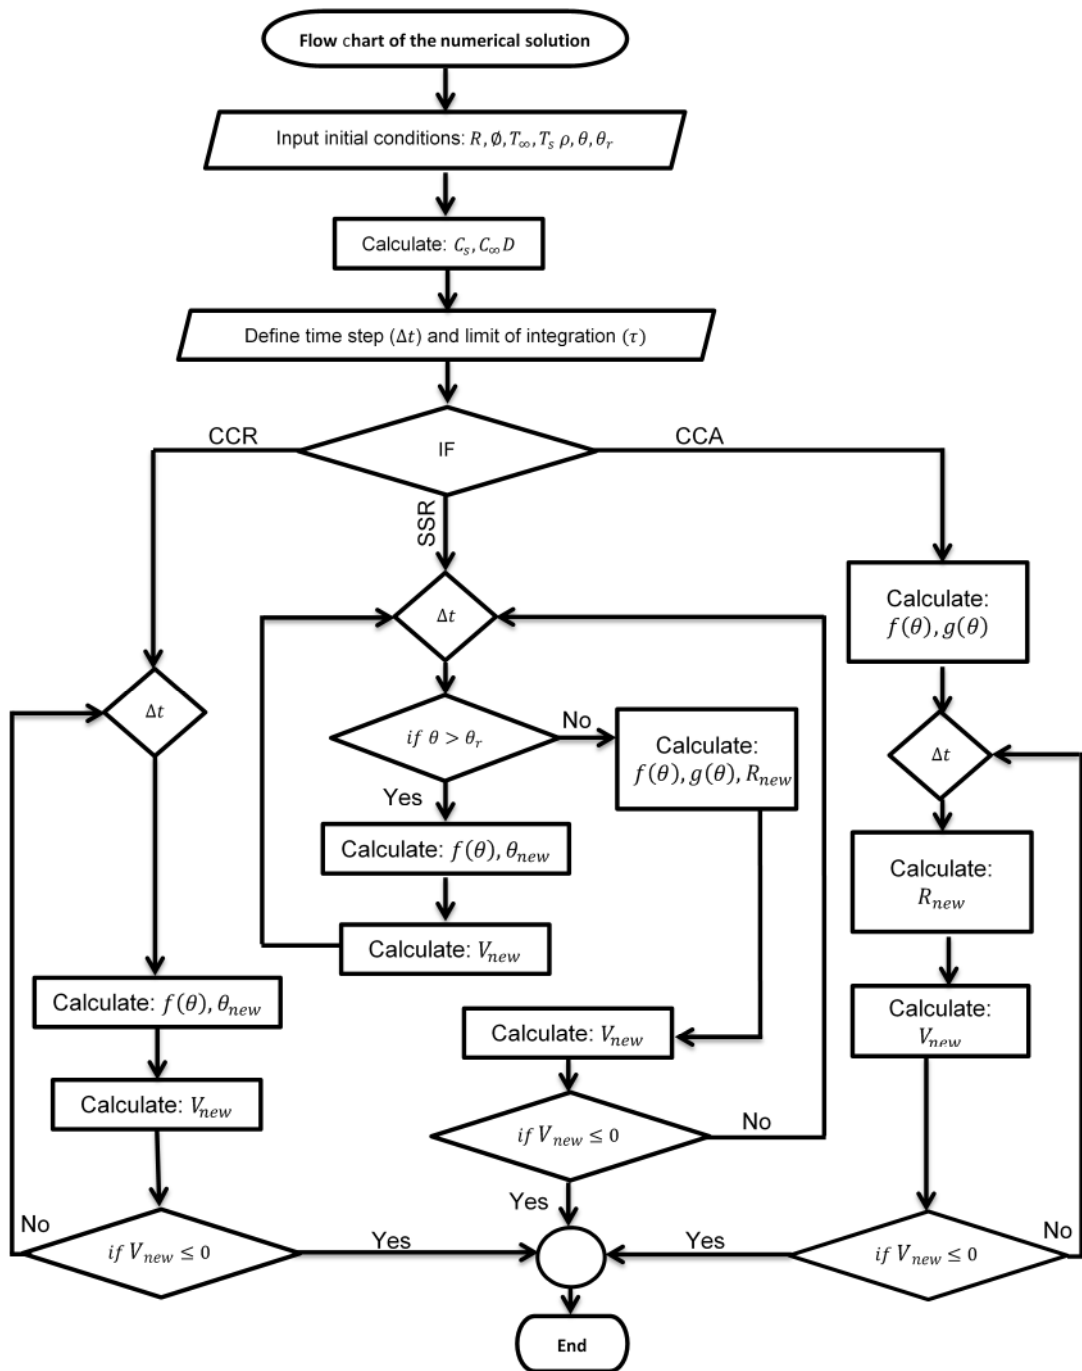
Figure 8.7 Calculation procedures of the diffusion only model based on the selected mode of evaporation.

8.5 Influence of the surface wettability and substrate temperature on the evaporation flux.

The spatial variation of vapour concentration based on the Deegan model, equation (8.5), is shown in figures 8.8 and 8.9 for the hydrophobic and hydrophilic surfaces, respectively. The arc length is calculated from the droplet surface radial, which is described by $\frac{r}{R} = \sinh\alpha/(\cosh\alpha + \cos\theta)$, and the vertical coordinate, which is described by $\frac{z}{R} = \sin\theta/(\cosh\alpha + \cos\theta)$. The local evaporation flux is presented for all the cases as a function of the normalized arc length which is defined as the line around the gas-liquid interface from the point of the gas-liquid-solid contact line to the point of maximum height. This is normalized by the total arc length of the droplet. Substrates temperatures of 40, 50 and 60℃ and initial droplet sizes of $2\mu l$ and $4\mu l$ are chosen to illustrate the difference in the variation of the evaporation flux on the droplet surface.

The evaporation flux at the start of the evaporation process, calculated using equation 8.5 for droplet sizes of $2\mu l$ and $4\mu l$ on the hydrophobic surface with different substrate temperatures is shown in figure 8.8. The results in figure 8.8 show that the evaporation flux is higher for the smaller droplet size; in all cases the maximum flux is at the point of maximum height of the droplet, and largely constant before rapidly reducing close to the contact line. This agrees with previous theoretical results [174], which show that the evaporation flux by natural diffusion (no substrate heating) of the sessile droplet with hydrophobic contact angle remains almost uniform along the upper half of the droplet surface. The behaviour with the hydrophilic surface is significantly different. Figure 8.9 shows equivalent data to that in figure 8.8, but for the initial conditions of the droplet with contact angle of 49°. It can be seen that the diffusion flux rapidly increases as the contact line is approached. In both cases reducing the temperature causes a reduction in evaporative flux, however the general overall pattern remains the same.
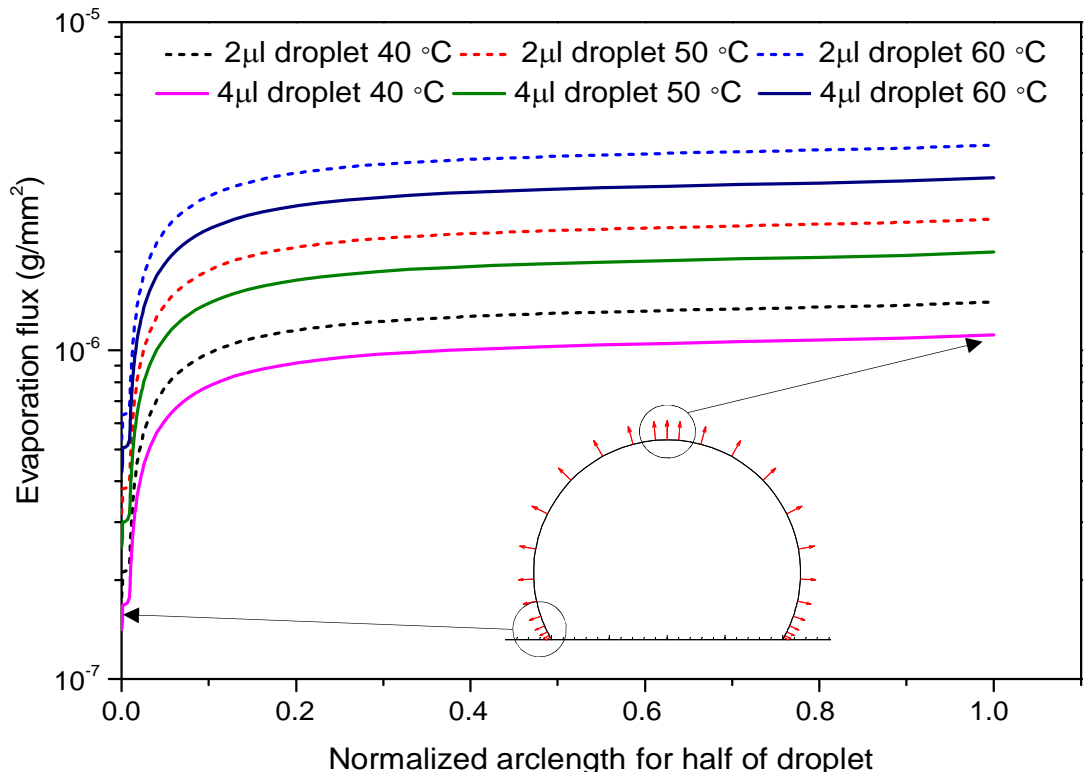
Figure 8.8 Local evaporation flux along 2µl and 4µl water droplet over aluminium hydrophobic surface for various substrate temperatures at time 0.
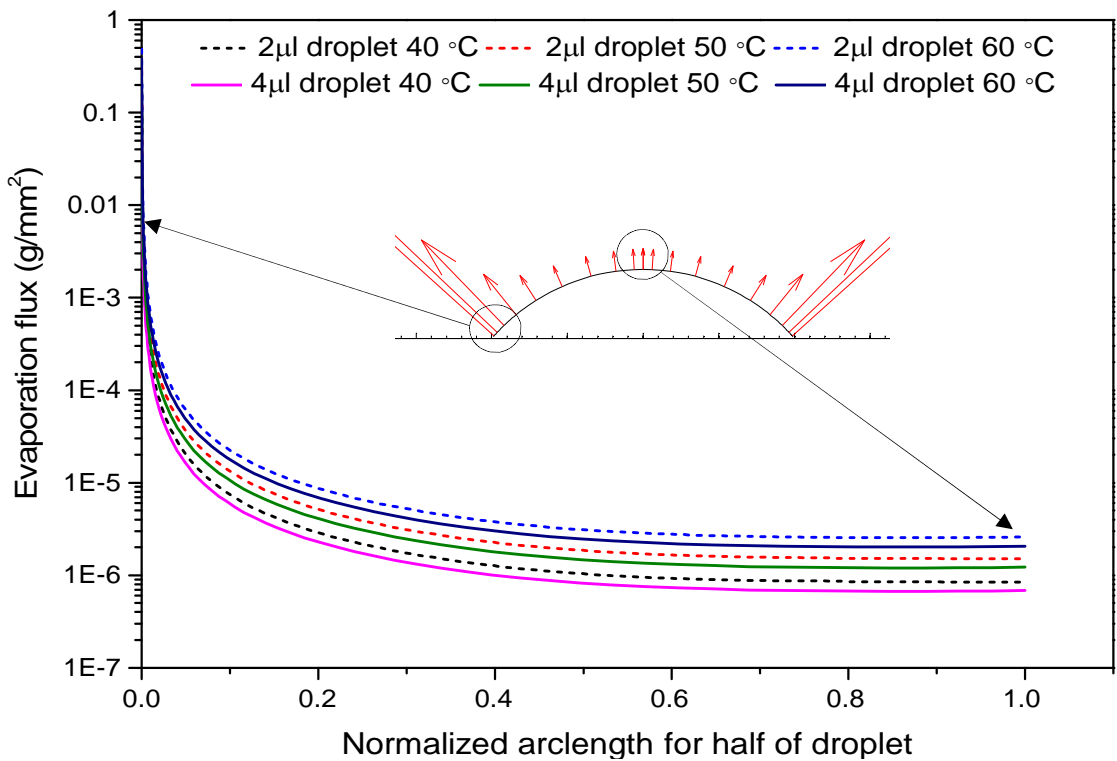
Figure 8.9 Local evaporation flux along 2µl and 4µl water droplet over aluminium hydrophilic surface for various substrate temperatures at time 0.

8.6 Total time of droplet evaporation over hydrophilic and hydrophobic surfaces

The total time of droplet evaporation over a surface is highly dependent on the surface temperature and wettability. Figures 8.10 shows the total time for the complete evaporation of 2, 4 and $8\mu l$ water droplets, for the hydrophobic and hydrophilic aluminium surfaces, together with the predicted values using the diffusion only model. Substrate temperature and wettability causes a reduction in the time taken for complete evaporation.

The stick-slip receding contact angle model (SSR) postulated here is considered for the mode of evaporation, where the droplet initially evaporates under a constant contact radius mode until the contact angle falls below the receding angle switching behaviour to a constant contact angle mode until complete evaporation. For most real surfaces, the modes of droplet evaporation are usually unknown a priori, however the combined model developed in this study (SSR) can mimic the physical behaviour of the droplet evaporation based on the measured static contact angle and the receding contact angle. Using this method will save a significant amount of time that is required to monitor the droplet life time for each surface required to be studied using the traditional visualization method to follow the receding behaviour when the droplet is de-pinning [171, 187, 188].

In the case of the aluminium hydrophobic surface for the droplet sizes dominated by the surface tension force (2, 4 and $8\,\mu l$), the total time of evaporation calculated by the vapour-diffusion model based on the SSR mode is generally in excellent agreement with that measured during the experiment where an average relative deviation of less than 4% is observed as shown in figures 8.10a-c. This means that the vapour-diffusion model, based on the stick-slip behaviour where the slip happens when the receding contact angle is reached, succeeds in describing the entire influence of substrate heating throughout the evaporation process. In the case of an aluminium hydrophilic surface, the results of the model are presented for the range of measured receding contact angles between 4 and 8°. However, a considerable under

237

prediction of 22% is observed when the surface is aluminium hydrophilic as shown in figures 8.10a-c.

The relative deviations between the experimental and analytical results are likely due to the disputing effect of the buoyant convection in the vapour phase [202] and the evaporative cooling along the interface [166] which are not included in the model. The buoyant convection is induced by the generated air flow around the droplet and substrate due to the temperature difference which enhances the evaporation rate of the droplet, while the evaporative cooling along the interface is induced mainly by the thermal resistance of the droplet layers and thermal conductivity of the substrate which reduces the evaporation rate of the droplet. The evaporative cooling along the interface is directly related to the ratio of the droplet height to the contact radius (h/R) while the buoyant convection was noted to increase with substrate temperature [162, 166, 202]. In the case of hydrophobic surface, the h/R ratio is high which gives an equivalent effect between the evaporative cooling across the interface and the buoyant convection resulting in good agreement between the experimental and analytical results for the Al-hydrophobic surface. On the other hand, the h/R aspect ratio is small for the hydrophilic surface resulting in minimizing the evaporative cooling along the interface. This supports the assertion for the mismatch between the measured and predicted total time of evaporation on the hydrophilic surface.
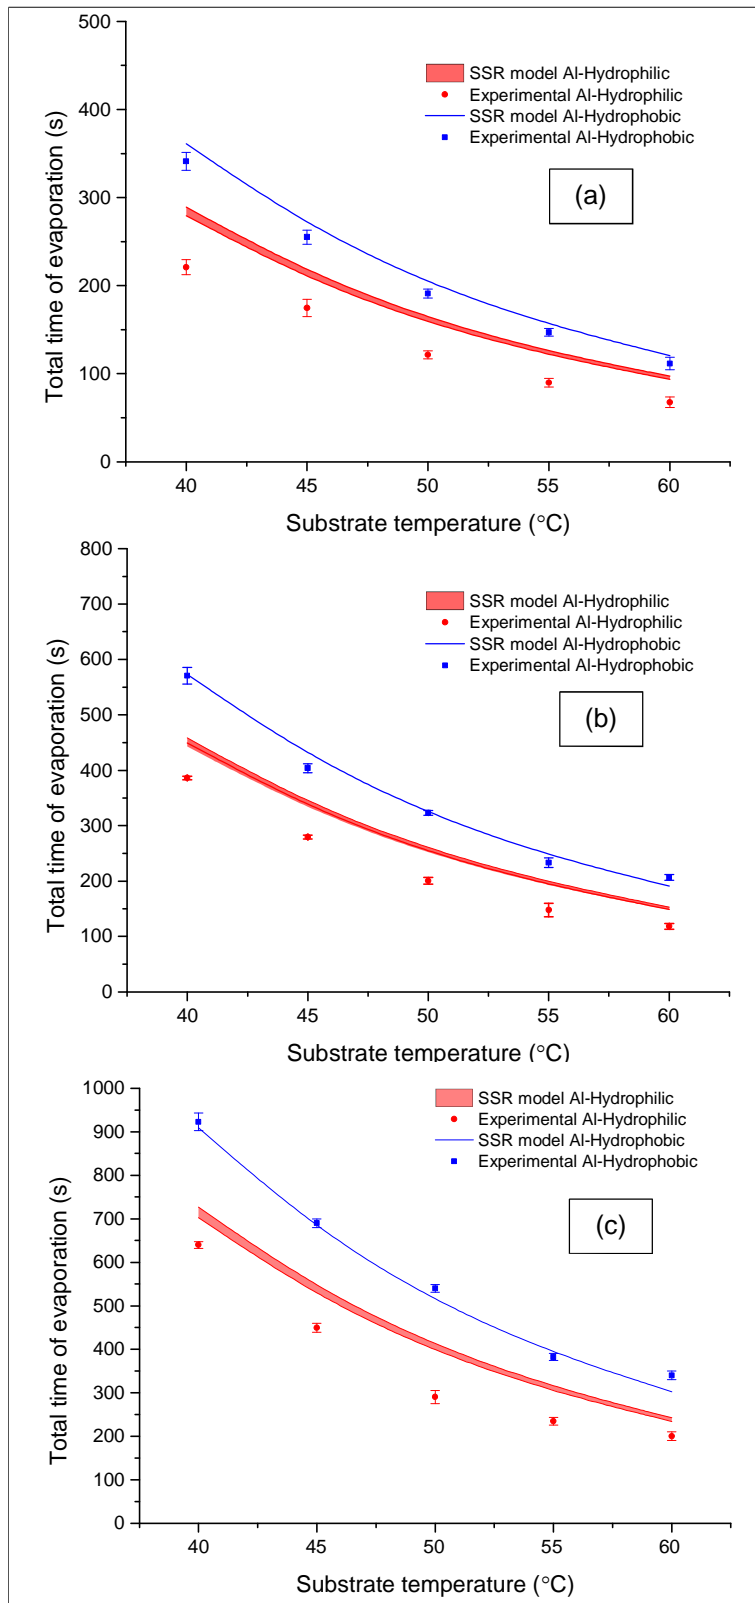
Figure 8.10 Total time for droplet evaporation on aluminium hydrophilic and hydrophobic surfaces as a function of substrate temperature. (a) 2µl, (b) 4µl and (c) 8µl.

8.7 Droplet weight evolution

8.7.1 Evaporation of droplet shapes dominated by surface tension

In this section, experimental and numerical results of the evaporation of water droplet are presented for Bond numbers in the range between 0.07 and 0.7 allowing the droplet shape to be considered as a spherical cap.

The reduction in droplet weight with time during evaporation of the droplets over the aluminium hydrophobic surface for three different temperatures $(40, 50 \text{ and } 60℃)$ are shown in figure 8.11 for droplet sizes of 2, 4 and 8μl. The evolution of weight is consistent with the earlier observation of droplet evaporation using the traditional visualization measurement technique over hydrophobic aluminium substrates coated with Teflon [166, 171] and PFC [162] coatings. The weight evolution of droplet evaporation over the aluminium hydrophobic surface was found to be predicted well by the vapour-diffusion model based on the stick-slip behaviour mode (SSR). This is in agreement with earlier observations using the CCA [166] and the CCR [162] modes.

The weight evolution of the water droplet over time during evaporation over the aluminium hydrophilic surface for three different temperatures $(40, 50 \text{ and } 60℃)$ is shown in figure 8.12 for the droplet sizes of $2, 4$ and $8\mu l$. As the initial contact angle is relatively small (49°) and the contact area is relatively large between the droplet and the surface, the temporal weight variation is rapidly decreasing with time, which is not the case for the hydrophobic surfaces. In the case of the hydrophilic surface, as mentioned previously, the vapour-diffusion model based on the SSR mode underpredicts the experimental evaporation rate with a percentage error of 22%. This mismatch between the experimental and theoretical results has been attributed to the high effect of the buoyant convection in the vapour phase compared to the evaporative cooling along the droplet interface which are not included in the present model.

The top and side views of an $8\mu l$ droplet on a heated substrate (60℃) is shown in table 8.1 for the initial droplet and at the point where 90% of mass has

evaporated. It is shown that the droplet shape can be measured easily using the traditional visualization methods at the initial stage. However, for the last 2% and 10% of the total time of evaporation of the hydrophobic and hydrophilic surfaces (respectively), there is a high associated error if one is calculating the droplet volume based on the contact angle and contact radius as the droplet shape is irregular. This is consistent with the previous observations [164, 191]. The high associated error in calculating the droplet size at the last stages of evaporation has resulted in a high uncertainty in calculating the total time of droplet evaporation and the rates of droplet evaporation are often extrapolated [154, 164, 192]. On the contrary to these results, the present way of weight measurements for the droplet evolution is shown to be superior regarding the independency on the droplet shape and the contact angle and can trace the evaporation process of any size of droplet which sits on any type of surface until the complete evaporation of the droplet without the need for any approximation or extrapolation.
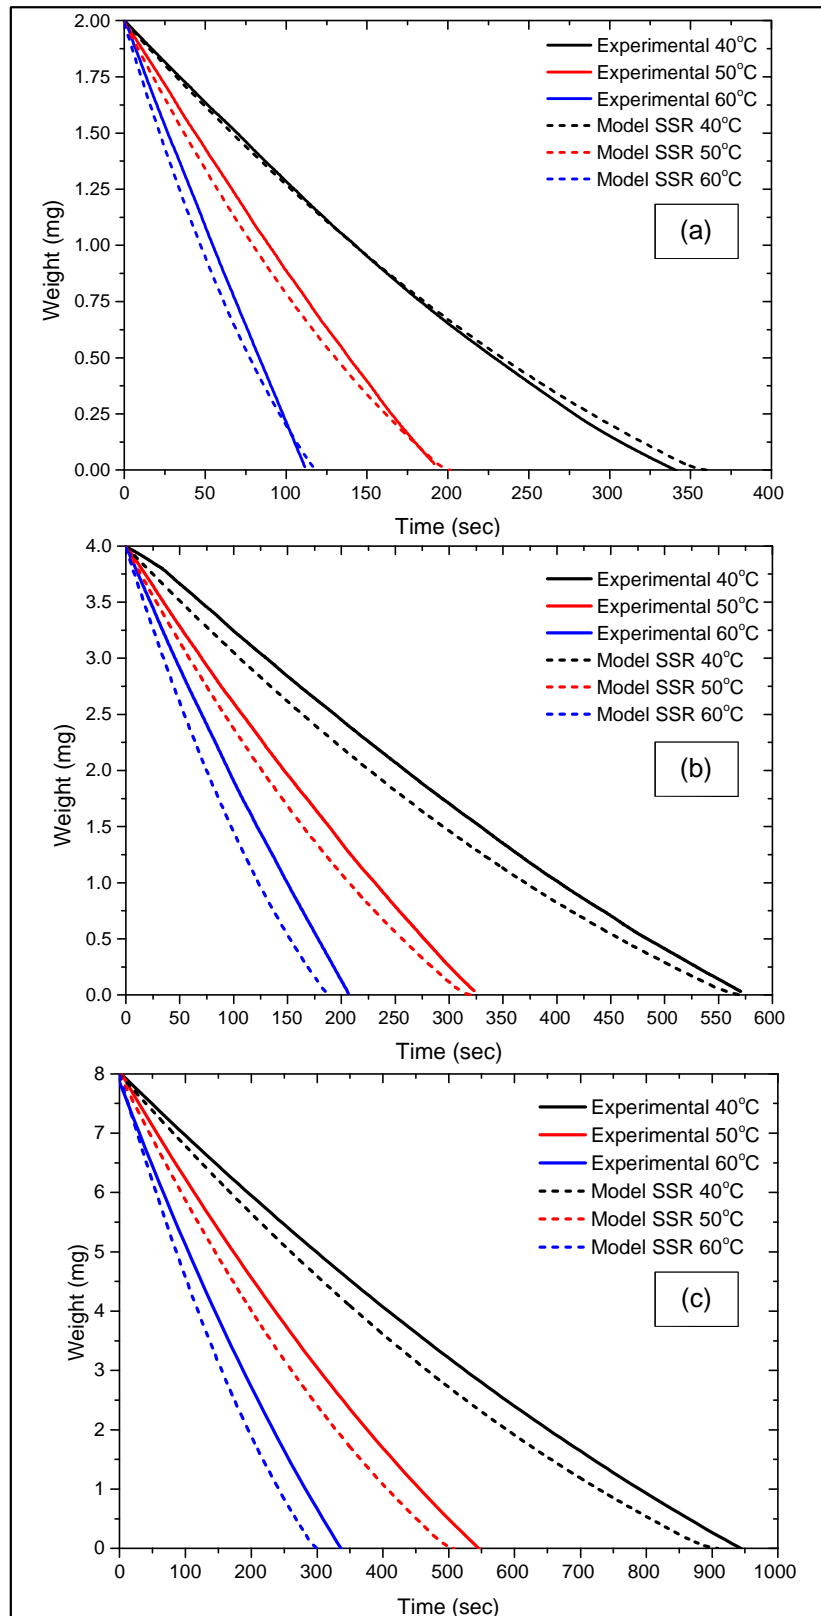
Figure 8.11 Temporal variation of droplet weight on the aluminium-hydrophobic surface. (a) 2µl, (b) 4µl and (c) 8µl.
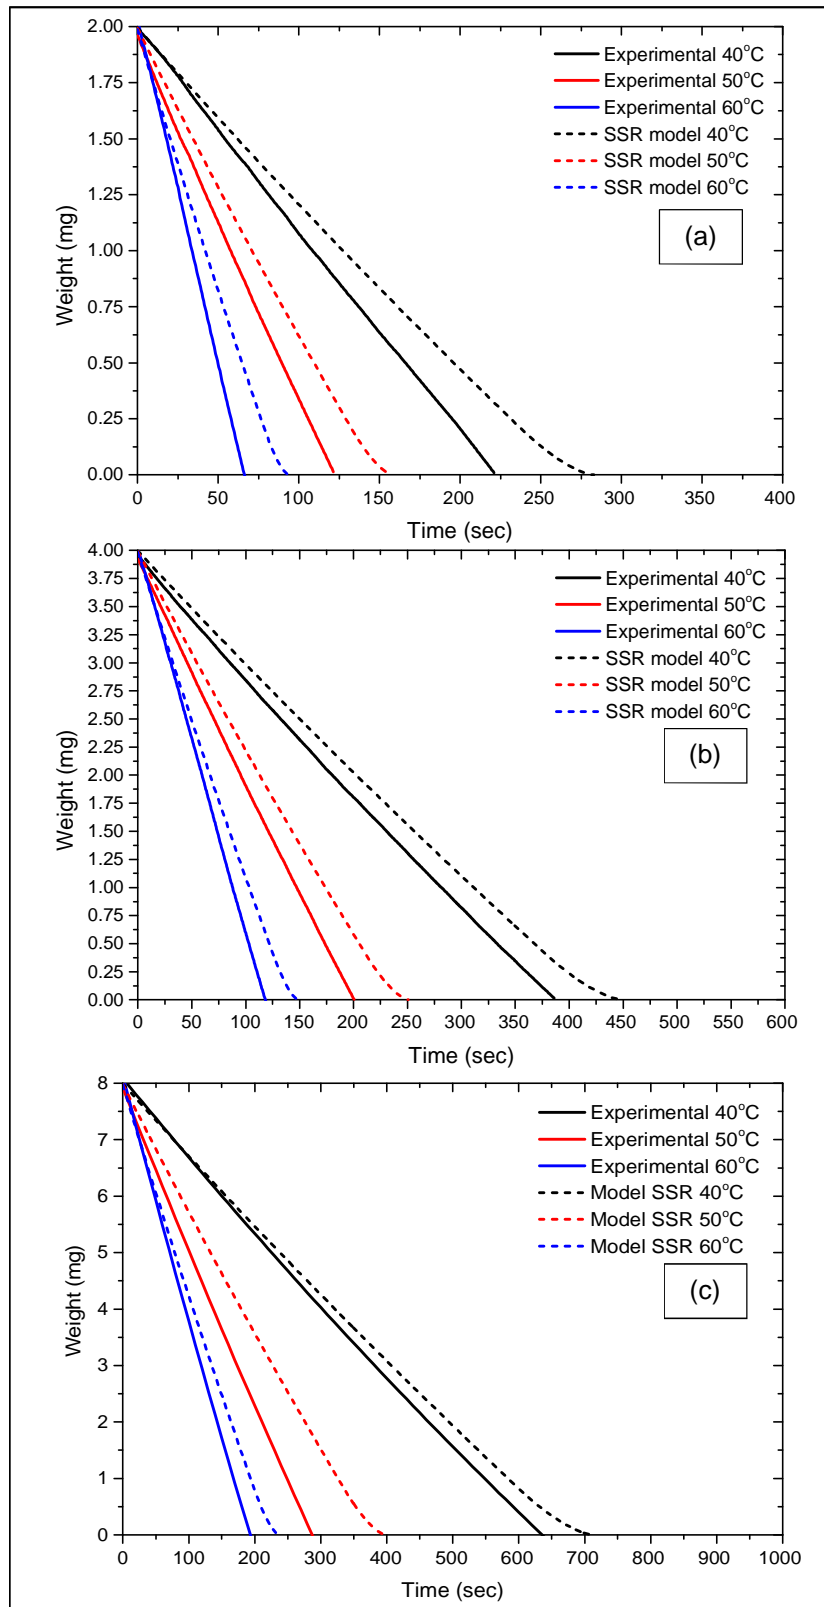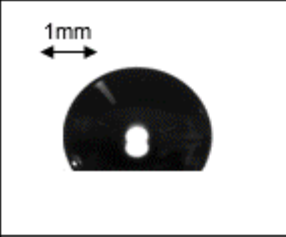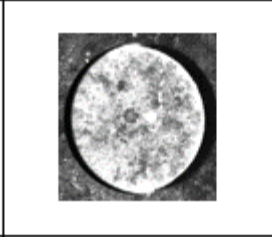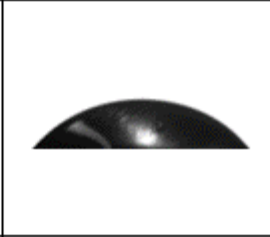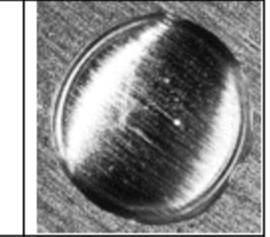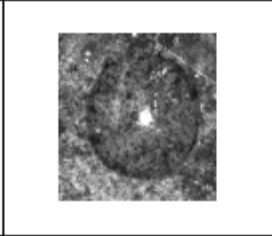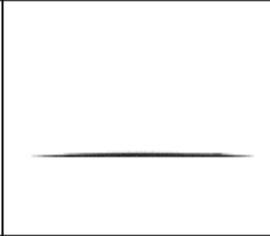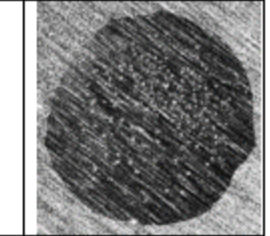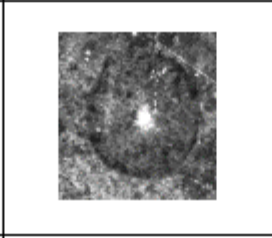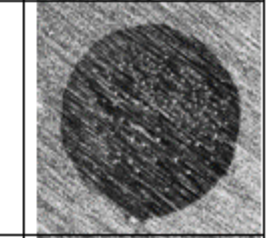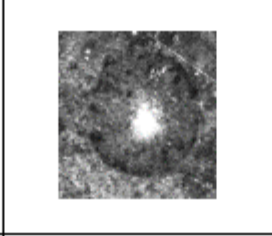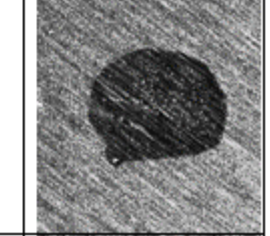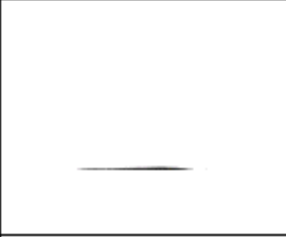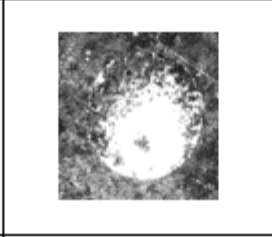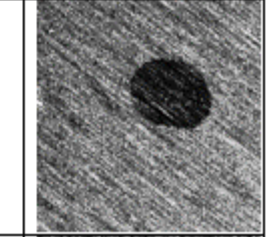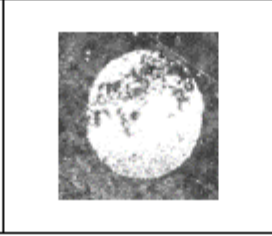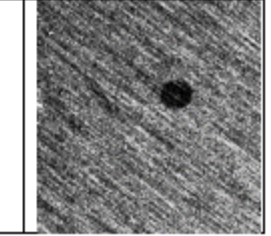
Figure 8.12 Temporal variation of droplet weight on the aluminium-hydrophilic surface. (a) 2μl, (b) 4μl and (c) 8μl.

Table 8.1 Time-dependent images of an 8µl water droplet during the evaporation under the substrate temperature of 60℃ for the Al-hydrophobic and Al-hydrophilic surfaces (t* is the normalized time which is the ratio between the instantaneous time and the total time of evaporation).

t^*	60℃-Al-Hydrophobic		60℃-Hydrophilic	
	Side view	Top view	Side view	Top view
0				
0.9				
0.92				
0.94				
0.96				
0.98				

1mm

8.7.2 Evaporation of droplet shapes dominated by gravity

In this section, the evaporation of relatively large water droplet volume ($30\mu l$) is studied. The temporal mass evolution of the droplet with time is shown in Figures 8.13a and 8.13b over the aluminium hydrophobic and hydrophilic surfaces respectively, for three different substrate temperatures (40, 50 and 60℃).

The chosen initial droplet volume gives higher gravity force over surface tension force which results in a Bond number larger than 1 (up to 1.7 for the hydrophilic surface). This is also associated with larger contact radius compared to the capillary length (k^{-1}) of the droplet ($k^{-1} = \sqrt{\frac{\gamma}{\rho g}}$). Therefore, the $30\mu l$ droplet shape is no longer considered as a spherical cap and the droplet surface tends to be flattened and contact area is irregular especially with the hydrophilic substrate. In both surfaces of hydrophobic and hydrophilic cases, the proposed method of measurement based on the weight evolution is successfully tracked by the droplet mass for the studied range of temperatures.

The SSR model was also used to predict the evaporation of the droplet for the range of temperatures and both types of surfaces as shown in figure 8.13. The percentage error in the hydrophobic case, shown in figure 8.13a, is 13% on average. However, the average error amplified to 47% for the hydrophilic case. The reason for the high percentage error is due to the irregularity in the droplet surface due to the high Bond number especially in the hydrophilic surface case and the assumption of the spherical cap of droplet in the model while it is intended to flatten with large initial droplet size.
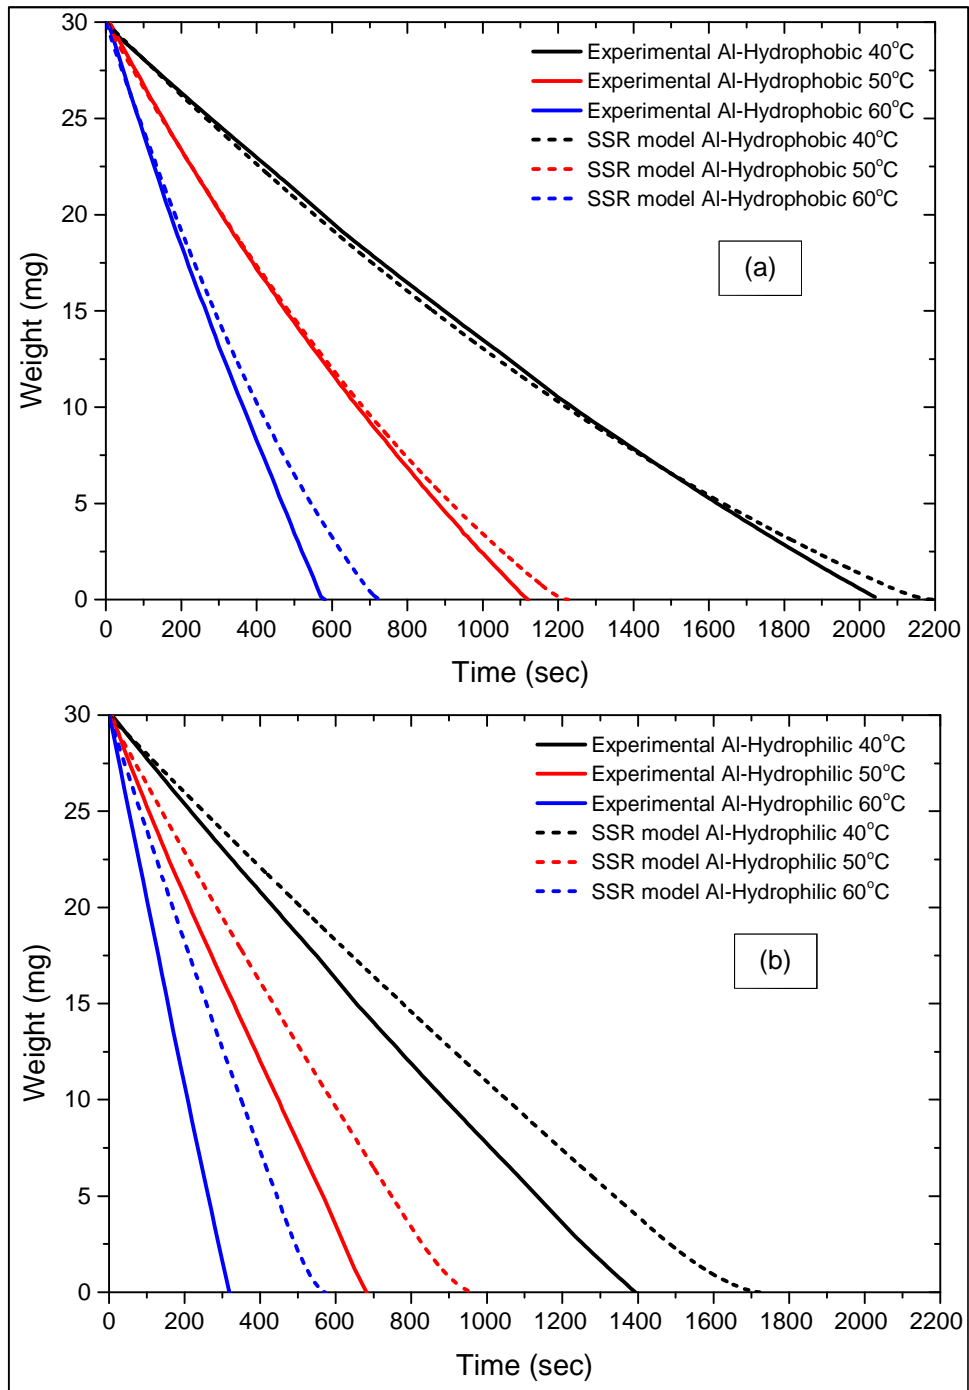
Figure 8.13 Temporal variation of 30µl droplet on the (a) aluminium hydrophobic and (b) aluminium hydrophilic surfaces.

8.8 Summary

A comprehensive experimental and theoritical investigation of droplet evaporation on heated hydrophilic and hydrophobic surfaces is provided. Measurements of the droplet evaporation is achieved using a precise analytical balance. A container to carry the PID controller circuit and the samples was designed to control the substrate temperature in the range of 40 to 60℃. The measurement technique has succefully managed to record and follow the total lifetime of any initial droplet size until complete evaporation.

A theoretical model is postulated here based on the stick-slip behaviour (SSR). To fully identify the point of shifting from CCR mode to CCA mode in the model, the receding contact angle is measured using the needle suction method to be considered as the point where the slip mode (CCA) starts. The theoretical predictions agree well with the experimental results for the aluminium hydrophobic surface with an average percentage error of 4%. The SSR model significantly underpredicts the evaporation rate for the aluminium hydrophilic surface with an average error of 22%. Moreover, the model showed that the error becomes significant for the big initial droplet sizes due to the increase in droplet irrgularities caused by the increase of the gravity force relative to the droplet surface tension force.

The evaporation flux was shown to be higher for smaller initial droplet size for both cases of hydrophobic and hydrophilic surfaces. This was also higher for droplet evaporating over hydrophilic surface compared with the hydrophobic surface which resulted in rapid evaporation of the droplet over the hydrophilic surface.

The result of this chapter has shown the importance of surface wettability to decrease the droplet evaporation time and increase the cooling performance. This cooncept can be appplied in heat exchangers that use spray evaporative cooling in data centres (AHU). However, this study can be extended in future research to cover the film evaporation over horizontal tubes and examine the effect of surface wetability which may be more pertinent than single droplet.

Chapter 9 Conclusion and recommendation

9.1 General discussion

The main aim of this thesis was to introduce better understanding of thermal and energy management of DCLC data centres through the integrated interactions between three different levels: server level, rack level and facility level. Therefore, multiple objectives were followed and achieved through the previous chapters.

In the first phase, a full scale of chiller-less DCLC data centre from the chip to the environment was designed. This included designing the control methodology, installations and commissioning. The main work involved two sections: IT environment and outdoor environment.

- The IT environment consisted of thirty Sun Fire V20z servers arranged in a standard data centre rack. These servers have their CPUs direct liquid cooled and the remaining of the components are air cooled. The heat generated by the CPUs of all the servers is collected through a coolant loop (secondary loop) and directly transferred to an external loop (primary loop) through a liquid/liquid heat exchanger identified as the CHx40. Whereas the heat generated by the air-cooled components of the servers is transferred to the external coolant loop using a passive rear door heat exchanger (RDHx).

- The outdoor environment involved a chiller-less heat rejection system which is called an Air Handling Unit (AHU) which represents the final point at which the heat is transferred to the environment. The AHU was designed to utilise spray evaporation to increase the performance of the heat exchanger. The primary loop is responsible for carrying the heat away from the IT environment to be transferred to the outdoor environment through the AHU.

The second phase in this thesis was to characterise and examine the designed cooling system to quantify the energy saving potentials and provide

the required means of cooling for safe operation of DCLC servers. This phase dealt with the whole facility (facility level) at which the dynamic responses of the components were characterised in terms of thermal and power consumption management. The system was also tested under long hour operation and for different IT computational workloads.

The third phase of the thesis involved focusing on the flow maldistribution in the rack and the resulted CPU temperature variations of different locations of the servers. The EPANET software was used to study and analyse the flow distribution validated by the experimental results. The flow distribution results were compared against the CPU temperatures of the servers. Moreover, the pumping configuration of the coolant in the rack was also investigated regarding the influences on the server component temperatures, power consumption and the effectiveness metrics (pPUE and TCUE) of the whole system.

The fourth phase constituted studying the effect of the coolant inlet temperature on the power consumption of the whole system, computational efficiency of the servers (WPE), power effectiveness metrics (pPUE and TCUE) and the relationship between the effectiveness and efficiency (pDWPE).

Finally, the last phase of this thesis involved investigation into the evaporation of droplets on surfaces. The underlying physics of this study can be applied in cold plate chip cooling and the AHU heat exchanger that utilise spray cooling. In this phase, the evaporation of sessile droplets on different heated surface wettability was investigated using a novel experimental technique based on weight measurements. Theoretical analysis based on the diffusion only model was developed to predict the droplet evaporation based on the initial experimental measurements of droplet size, static contact angle and the receding contact angle using the needle suction method. By using this method, it can easily predict the receding action of the droplet without the need for full imaging experiment.

9.2 Research contribution

The scope of this work included design of a chiller-less liquid cooled data centre and evaluated holistically the impact of ambient conditions, workload, flow, and temperature on the thermal, energy and IT computational performance. This investigation helps to introduce a better understanding of the integrated interactions between different levels of DCLC data centres. The contributions and novelty of the current work are summarised as follow:

1. Presented full design guidelines for chiller-less thermal and energy efficient direct contact liquid cooled data centre.

2. Identified the temperature variation of the CPUs of different servers in the rack and determined the causes of these differences.

3. Investigated the flow maldistribution in the DCLC rack and suggested recommendation for running conditions, operation and design to weaken and eliminate the flow maldistribution.

4. Evaluated the impact of different methods of pumping the coolant in the rack on the thermal and energy performance.

5. Evaluated the IT inlet coolant temperature on the thermal, energy and computational performance of the DCLC rack.

6. Considered the data centre holistically (DWPE) and combined the compute efficiency (FLOPS/W) and power effectiveness (PUE).

7. Introduced a new technique to measure droplet evaporation on heated surfaces for any droplet surface tension force.

8. Introduced a new method of simulating droplet evaporation based on the initial conditions of static contact angles and receding contact angles regardless of the evaporation mode.

9.3 Concluding remarks

As mentioned previously, the aim of this thesis was achieved through multiple objectives. Therefore, the findings of this thesis will be summarised and categorised based on these objectives as follows:

I.   The main conclusion points of the design of the chiller-less DCLC cooling system and the characterisation of the components and operation are as follows:

1. The designed (bespoke) heat rejection system can maintain a constant supplied cooling temperature to the servers despite variations in environmental and digital workloads.

2. Using DCLC with an AHU and an evaporative spray system has the potential to operate at a low pPUE/TCUE.

3. Enabling the spray should be tightly controlled to make sure that the heat rejection system is not consuming more power than necessary – resulting in a higher pPUE/TCUE.

4. Using the spray evaporation found to reduce the thermal resistance of the heat exchanger by 50%. However, only 25% of the spray flow rate was found to evaporate and being converted into useful latent heat of evaporation.

5. High pressure drop is associated with higher power consumption and less thermal benefits. The pressure drop across the heat exchanger of the AHU was found to increase by 30% when increasing the fan speed from 250rpm to 1000rpm. The fan should be chosen properly to overcome the pressure drop especially when using the spray in the design as adding the spray increase the pressure drop up to 145% across the heat exchanger.

6. A lag of 10 min between the actions producing more heat in the servers and the response of the AHU was recorded. Therefore, it is important to consider this time lag in future designs.

7. CPU0 of a server in the rack generally have a lower temperature than CPU1. This is justified by the fact that CPU1 receives higher coolant

temperature than CPU0. However, depending on the IT load distribution, CPU1 can have a lower temperature than CPU0 in some cases.

II. The main concluding remarks of the flow distribution and flow configuration inside the rack are:

1. The EPANET software model was able to predict the flow rate in the rack with a good accuracy of 2.3% compared with the experimental results.

2. The current design showed that the server in top of the rack (`eng01`) receives a greater coolant flow rate than the server in the bottom of the rack (`eng30`), namely an increase of 28% and 30% for the centralised and distributed pumping configurations, respectively.

3. The flow maldistribution resulted in a general increase in the CPU temperatures of the servers from the top to the bottom of the rack.

4. The flow maldistribution in the servers is due mainly to two factors: friction losses in the manifolds, and the pressure drop variations in the server modules.

5. The optimisation analysis showed that increasing the hydraulic diameter of the manifolds from 2.5 to 5cm yields an enhancement in the flow distribution to the servers. Moreover, the overall flow rate was also found to increase by 9% and 10% for the central and distributed pumping configurations, respectively.

6. The CPU temperatures were shown to be higher by 2℃ in the distributed pumping configuration case compared with the central pumping configuration for high IT loads for the same flow rate.

7. The CPU temperatures are clearly affected by the total flow rate of the coolant in the rack at the high IT load levels.

8. Some servers encountered shutdown at low coolant flow rates in the central pumping case. Therefore, it can be very effective to control the server module pumps to switch on using an algorithm based on the CPU temperature of the server.

9. The pPUE and TCUE deteriorate with higher flow rate as more power is added to the cooling side. However, the TCUE is more sensitive to the IT load as the servers' fans ramp up, which adds more power to the cooling side while this increase in power is added to the IT side for the pPUE.

III. The main conclusion points for the influence of the rack inlet coolant temperature are as follows:

1. The CPU temperatures are profoundly linked to the rack inlet temperature and shutdown is encountered for a couple of servers when the rack inlet temperature was higher than 37℃ for the IT loads higher than 50%.

2. Increasing the rack inlet temperature found to reduce the AHU power consumption but also increase the rack power consumption. Therefore, both of these effects have helped to lower the pPUE and TCUE metrics.

3. Depending on the IT load level, supplying higher coolant inlet temperature to the rack reduces the computational performance of the rack. For example, at 75% IT load level, the WPE decreased by 4% when the rack inlet temperature increased from 22 to 44℃.

4. This work highlighted the important fact that high effectiveness of the data centre does not necessarily mean high efficiency. As the pPUE and TCUE could be improved by supplying less power for the cooling system while the servers performance can deteriorate.

5. Considering the effectiveness and efficiency of the current rack design of the Sun Fire V20z cloud servers, supplying an inlet temperature of 30℃ can be ideal to work in the safe operation conditions.

IV. The most important conclusions for the droplet evaporation on heated surfaces are:

1. The proposed weight measurement technique of droplet evaporation over heated surfaces has successfully shown to record the whole life time of the droplet enabling study of any size of droplet from surface tension domination to gravity domination.

2. The evaporation behaviour and total time of evaporation are strongly linked to the surface wettability, initial droplet size and substrate temperature. The evaporation rate improves considerably with lower contact angle, smaller initial droplet size and higher substrate temperature.

3. The proposed SSR model predicts the total time of evaporation over hydrophobic surfaces to good accuracy with less than 4% difference between theory and experiment. However, the model showed a significant underestimation of the total evaporation rate with an error of 22% for the hydrophilic surfaces.

9.4 Recommendation for future work

The following points highlight some aspects which are suggested based on this thesis to be covered as possible future work:

1. Investigate the use of chiller-less DCLC design on various servers from different venders.

2. Evaluate the performance of the servers under real data centre operations.

3. Extend the study to different environmental conditions to cover the performance of the AHU for a wide range of locations.

4. Extend the study to cover the dynamic response and failure scenarios that might occur in different components of the designed cooling system.

5. Develop an analytical model to describe the heat transfer from the chip to the environment in the DCLC chiller-less arrangement.

6. Study the feasibility of using different heat exchanger inclination in the AHU with spray evaporation on the cooling performance.

7. Extend the study to cover the behaviour of the troublesome servers by enabling their module pumps while other module pumps are turned off.

8. Develop an algorithm to control the server module pumps operation based on the need and the component temperature of the servers and study the feasibility of this approach.

9. Apply the possibility of using droplet evaporation on real cold plate surfaces to explore the potential of increasing the heat transfer capacity.

10. Study the effect of the AHU heat exchanger surface wettability on the cooling performance under spray evaporation.

References

1. Joshi, Y. and P. Kumar, *Energy efficient thermal management of data centers.* 2012: Springer Science & Business Media.

2. Khalaj, A.H. 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, 2017. **205**: p. 1165-1188.

3. Patankar, S.V., *Airflow and cooling in a data center.* Journal of Heat Transfer, 2010. **132**(7): p. 073001.

4. Rahmat Romadhon, Maisarah Ali, Ayub Mohd Mahdzir, and Y.A. Abakr, *Optimization of Cooling Systems in Data Centre by Computational Fluid Dynamics Model and Simulation*, in *Conference on Innovative Technologies in Intelligent Systems and Industrial Applications (CITISIA 2009).* 2009: Monash University, Sunway campus, Malaysia,. p. 322-327.

5. Mitchell-Jackson, J., J.G. Koomey, B. Nordman, and M. Blazek, *Data center power requirements: measurements from Silicon Valley.* Energy, 2003. **28**(8): p. 837-850.

6. Schultz, M., M. Gaynes, P. Parida, and T. Chainer. *Experimental investigation of direct attach microprocessors in a liquid-cooled chiller-less data center.* in *Thermal and Thermomechanical Phenomena in Electronic Systems (ITherm), 2014 IEEE Intersociety Conference on.* 2014. IEEE.

7. Brown, R.E., R. Brown, E. Masanet, B. Nordman, B. Tschudi, A. Shehabi, J. Stanley, J. Koomey, D. Sartor, and P. Chan, *Report to congress on server and data center energy efficiency: Public law 109-431.* 2007, Ernest Orlando Lawrence Berkeley National Laboratory, Berkeley, CA (US).

8. Koomey, J., *Growth in data center electricity use 2005 to 2010.* A report by Analytical Press, completed at the request of The New York Times, 2011.

9. Ebrahimi, K., 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, 2014. **31**: p. 622-638.

10. Koomey, J.G., *Worldwide electricity used in data centers.* Environmental research letters, 2008. **3**(3): p. 034008.

11. Garimella, S.V., T. Persoons, J. Weibel, and L.-T. Yeh, *Technological drivers in data centers and telecom systems: Multiscale thermal, electrical, and energy management.* Applied energy, 2013. **107**: p. 66-80.

12. VanGeet, O., *Best practices guide for energy-efficient data center design.* 2010, EERE Publication and Product Library.

13. Alkharabsheh, S., J. Fernandes, B. Gebrehiwot, D. Agonafer, K. Ghose, A. Ortega, Y. Joshi, and B. Sammakia, *A Brief Overview of Recent Developments in Thermal Management in Data Centers.* Journal of Electronic Packaging, 2015. **137**(4): p. 040801.

14. Rambo, J.D., *Reduced-order modeling of multiscale turbulent convection: application to data center thermal management.* 2006, Georgia Institute of Technology.

15. ASHRAE, T., *Datacom equipment power trends and cooling applications.* 2012, American Society of Heating, Refrigeration and Air-Conditioning Engineers Atlanta, GA.

16. Kheirabadi, A.C. and D. Groulx, *Cooling of server electronics: A design review of existing technology.* Applied Thermal Engineering, 2016.

17. Chu, R.C., R.E. Simons, M.J. Ellsworth, R.R. Schmidt, and V. Cozzolino, *Review of cooling technologies for computer products.* IEEE Transactions on Device and materials Reliability, 2004. **4**(4): p. 568-585.

18. Ohadi, M., S. Dessiatoun, K. Choo, M. Pecht, and J.V. Lawler. *A comparison analysis of air, liquid, and two-phase cooling of data centers.* in *Semiconductor Thermal Measurement and Management Symposium (SEMI-THERM), 2012 28th Annual IEEE.* 2012. IEEE.

19. Steinbrecher, R.A., *Data Center Environments ASHRAE's Evolving Thermal Guidelines.* ASHRAE Journal, 2011. **53**(12): p. 42.

20. ASHRAE, *Thermal Guidelines for Data Processing Environments*, ed. T. Edition. 2012. 152.

21. Jamalzadeh, M. and N. Behravan, *An exhaustive framework for better data centers' energy efficiency and greenness by using metrics.* Indian Journal of Computer Science and Engineering, 2012. **2**(6): p. 813-822.

22. Mathew, P., *Self-benchmarking guide for data centers: Metrics, benchmarks, actions.* Lawrence Berkeley National Laboratory, 2010.

23. Sharma, R.K., C.E. Bash, and C.D. Patel. *Dimensionless parameters for evaluation of thermal design and performance of large-scale data centers*. in *8th ASME/AIAA Joint Thermophysics and Heat Transfer Conference*. 2002.

24. Herrlin, M.K., *Rack cooling effectiveness in data centers and telecom central offices: The rack cooling index (RCI).* Transactions-American Society of Heating Refrigerating and Air conditioning Engineers, 2005. **111**(2): p. 725.

25. Herrlin, M.K. *Improved data center energy efficiency and thermal performance by advanced airflow analysis*. in *Digital Power Forum*. 2007.

26. Vangilder, J.W. and S.K. Shrivastava, *Capture Index: An Airflow-Based Rack Cooling Performance Metric.* ASHRAE transactions, 2007. **113**(1).

27. Lettieri, D.J., C.R. Hannemann, V.P. Carey, and A.J. Shah. *Lifetime exergy consumption as a sustainability metric for information technologies*. in *Sustainable Systems and Technology, 2009. ISSST'09. IEEE International Symposium on*. 2009. IEEE.

28. Brady, G.A., *Energy Efficiency in Data Centres and the Barriers to Further Improvements: An Interdisciplinary Investigation*. 2016, University of Leeds.

29. Malone, C. and C. Belady. *Metrics to characterize data center & IT equipment energy use*. in *Proceedings of the Digital Power Forum, Richardson, TX*. 2006.

30. Consortium, G.G.I., *Green Grid Metrics Describing Data Center Power Efficiency.* Technical Committee White Paper, available at http://www. thegreengrid. org/~/media/WhitePapers/Green_Grid_Metrics_WP. pdf, 2007.

31. Stansberry, M., *uptime institute annual data center industry survey report and full results.".* 2013.

32. Brady, G.A., N. Kapur, J.L. Summers, and H.M. Thompson, *A case study and critical assessment in calculating power usage effectiveness for a data centre.* Energy Conversion and Management, 2013. **76**: p. 155-161.

33. Horner, N. and I. Azevedo, *Power usage effectiveness in data centers: overloaded and underachieving.* The Electricity Journal, 2016. **29**(4): p. 61-69.

34. *BS EN 50600-4-2:2016: Information technology. Data centre facilities and infrastructures. Power Usage Effectiveness*. 2017, British Standards Institute.

35. Sun, H. and S. Lee, *Case study of data centers' energy performance.* Energy and Buildings, 2006. **38**(5): p. 522-533.

36. Tschudi, B., O. Vangeet, J. Cooley, and D. Azevedo, *ERE: A metric for measuring the benefit of reuse energy from a data center.* White Paper, 2010. **29**.

37. Shah, A.J., C.E. Bash, N. Kumari, T. Cader, T.J. Breen, E.J. Walsh, and J. Punch. *On the need for energy efficiency metrics that span integrated IT-facility infrastructures.* in *ASME 2011 Pacific Rim Technical Conference and Exhibition on Packaging and Integration of Electronic and Photonic Systems.* 2011. American Society of Mechanical Engineers.

38. Patterson, M.K., S.W. Poole, C.-H. Hsu, D. Maxwell, W. Tschudi, H. Coles, D.J. Martinez, and N. Bates. *TUE, a new energy-efficiency metric applied at ORNL's Jaguar.* in *International Supercomputing Conference.* 2013. Springer.

39. Green500. Available from: https://www.Green500.org/.

40. Minartz, T., J.M. Kunkel, and T. Ludwig, *Simulation of power consumption of energy efficient cluster hardware.* Computer Science-Research and Development, 2010. **25**(3-4): p. 165-175.

41. Newcombe, L., Z. Limbuwala, P. Latham, and V. Smith, *Data centre fixed to variable energy ratio metric (dc-fver).* British Computer Society, 2012.

42. Subramaniam, B. and W.-c. Feng. *The Green Index: A metric for evaluating system-wide energy efficiency in HPC systems.* in *Parallel and Distributed Processing Symposium Workshops & PhD Forum (IPDPSW), 2012 IEEE 26th International.* 2012. IEEE.

43. Beitelmal, A. and D. Fabris, *Servers and data centers energy performance metrics.* Energy and Buildings, 2014. **80**: p. 562-569.

44. Wilde, T., A. Auweter, M.K. Patterson, H. Shoukourian, H. Huber, A. Bode, D. Labrenz, and C. Cavazzoni. *DWPE, a new data center energy-efficiency metric bridging the gap between infrastructure and workload.* in *High Performance Computing & Simulation (HPCS), 2014 International Conference on.* 2014. IEEE.

45. Sharma, R., A. Shah, C. Bash, T. Christian, and C. Patel. *Water efficiency management in datacenters: Metrics and methodology*. in *Sustainable Systems and Technology, 2009. ISSST'09. IEEE International Symposium on*. 2009. IEEE.

46. Azevedo, D., M. Patterson, J. Pouchet, and R. Tipley, *Carbon usage effectiveness (CUE): a green grid data center sustainability metric.* White Paper, 2010. **32**.

47. Li, Z. and S.G. Kandlikar, *Current status and future trends in data-center cooling technologies.* Heat Transfer Engineering, 2015. **36**(6): p. 523-538.

48. Geng, H., *Data center handbook*. 2014: John Wiley & Sons.

49. Chiriac, V.A. and F. Chiriac. *Novel energy recovery systems for the efficient cooling of data centers using absorption chillers and renewable energy resources*. in *Thermal and Thermomechanical Phenomena in Electronic Systems (ITherm), 2012 13th IEEE Intersociety Conference on*. 2012. IEEE.

50. Fakhim, B., M. Behnia, S. Armfield, and N. Srinarayana, *Cooling solutions in an operational data centre: A case study.* Applied thermal engineering, 2011. **31**(14): p. 2279-2291.

51. Oró, E., V. Depoorter, A. Garcia, and J. Salom, *Energy efficiency and renewable energy integration in data centres. Strategies and modelling review.* Renewable and Sustainable Energy Reviews, 2015. **42**: p. 429-445.

52. Committee, A.T., *Thermal guidelines for data processing environments-expanded data center classes and usage guidance.* Atlanta, 2011.

53. Xu, G., B. Guenin, and M. Vogel. *Extension of air cooling for high power processors*. in *Thermal and Thermomechanical Phenomena in Electronic Systems, 2004. ITHERM'04. The Ninth Intersociety Conference on*. 2004. IEEE.

54. Katoh, T., G. Xu, M. Vogel, and S. Novotny. *New attempt of forced-air cooling for high heat-flux applications*. in *Thermal and Thermomechanical Phenomena in Electronic Systems, 2004. ITHERM'04. The Ninth Intersociety Conference on*. 2004. IEEE.

55. Zhang, H., S. Shao, H. Xu, H. Zou, and C. Tian, *Free cooling of data centers: A review.* Renewable and Sustainable Energy Reviews, 2014. **35**: p. 171-182.

56. Patankar, S.V. and K.C. Karki, *Distribution of Cooling Airflow in a Raised-Floor Data Center.* ASHRAE Transactions, 2004. **110**(2).

57. Muralidharan, B., M. Ibrahim, S. Shrivistava, S. Alkharabsheh, and B. Sammakia, *Impact of cold aisle containment on thermal performance of data center.* ASME Paper No. IPACK2013-73201, 2013.

58. Garimella, S.V., L.-T. Yeh, and T. Persoons, *Thermal management challenges in telecommunication systems and data centers.* IEEE Transactions on Components, Packaging and Manufacturing Technology, 2012. **2**(8): p. 1307-1316.

59. Srinarayana, N., B. Fakhim, M. Behnia, and S.W. Armfield, *Thermal performance of an air-cooled data center with raised-floor and non-raised-floor configurations.* Heat Transfer Engineering, 2014. **35**(4): p. 384-397.

60. Walsh, E.J., T.J. Breen, J. Punch, A.J. Shah, and C.E. Bash. *From chip to cooling tower data center modeling: Part II influence of chip temperature control philosophy.* in *Thermal and Thermomechanical Phenomena in Electronic Systems (ITherm), 2010 12th IEEE Intersociety Conference on*. 2010. IEEE.

61. Tuma, P.E. *The merits of open bath immersion cooling of datacom equipment.* in *Semiconductor Thermal Measurement and Management Symposium, 2010. SEMI-THERM 2010. 26th Annual IEEE*. 2010. IEEE.

62. Tuma, P.E. *Design considerations relating to non-thermal aspects of passive 2-phase immersion cooling.* in *Semiconductor Thermal Measurement and Management Symposium (SEMI-THERM), 2011 27th Annual IEEE*. 2011. IEEE.

63. El-Genk, M.S., *Immersion cooling nucleate boiling of high power computer chips.* Energy conversion and management, 2012. **53**(1): p. 205-218.

64. Asadi, M., G. Xie, and B. Sunden, *A review of heat transfer and pressure drop characteristics of single and two-phase microchannels.* International Journal of Heat and Mass Transfer, 2014. **79**: p. 34-53.

65. Kadam, S.T. and R. Kumar, *Twenty first century cooling solution: Microchannel heat sinks.* International Journal of Thermal Sciences, 2014. **85**: p. 73-92.

66. Al-Neama, A.F., N. Kapur, J. Summers, and H.M. Thompson, *An experimental and numerical investigation of the use of liquid flow in serpentine microchannels for microelectronics cooling.* Applied Thermal Engineering, 2017. **116**: p. 709-723.

67. Ellsworth, M.J., G.F. Goth, R.J. Zoodsma, A. Arvelo, L.A. Campbell, and W.J. Anderl, *An overview of the IBM power 775 supercomputer water cooling system.* Journal of Electronic Packaging, 2012. **134**(2): p. 020906.

68. Burd, T.D. and R.W. Brodersen, *Processor design for portable systems*, in *Technologies for wireless computing*. 1996, Springer. p. 119-137.

69. Sarood, O., A. Gupta, and L.V. Kale. *Temperature aware load balancing for parallel applications: Preliminary work*. in *Parallel and distributed processing workshops and phd forum (IPDPSW), 2011 IEEE international symposium on*. 2011. IEEE.

70. Menon, H., B. Acun, S.G. De Gonzalo, O. Sarood, and L. Kalé. *Thermal aware automated load balancing for hpc applications*. in *Cluster Computing (CLUSTER), 2013 IEEE International Conference on*. 2013. IEEE.

71. Ge, R., X. Feng, and K.W. Cameron. *Performance-constrained distributed dvs scheduling for scientific applications on power-aware clusters*. in *Supercomputing, 2005. Proceedings of the ACM/IEEE SC 2005 Conference*. 2005. IEEE.

72. Meijer, G.I., *Cooling energy-hungry data centers.* Science, 2010. **328**(5976): p. 318-319.

73. Kandlikar, S.G. and C.N. Hayner, *Liquid cooled cold plates for industrial high-power electronic devices—Thermal design and manufacturing considerations.* Heat Transfer Engineering, 2009. **30**(12): p. 918-930.

74. Ouchi, M., Y. Abe, M. Fukagaya, H. Ohta, Y. Shinmoto, M. Sato, and K.-i. Iimura. *Thermal management systems for data centers with liquid cooling technique of CPU*. in *Thermal and Thermomechanical Phenomena in Electronic Systems (ITherm), 2012 13th IEEE Intersociety Conference on*. 2012. IEEE.

75. Goth, G.F., A. Arvelo, J. Eagle, M.J. Ellsworth, K.C. Marston, A.K. Sinha, and J.A. Zitz. *Thermal and mechanical analysis and design of the IBM Power 775 water cooled supercomputing central electronics complex*. in *Thermal and Thermomechanical Phenomena in Electronic Systems (ITherm), 2012 13th IEEE Intersociety Conference on*. 2012. IEEE.

76. Kheirabadi, A.C. and D. Groulx, *Experimental Evaluation of a Thermal Contact Water Cooling System for Server Electronics* 2017.

77. Ramakrishnan, B., S. Alkharabsheh, Y. Hadad, B. Sammakia, P.R. Chiarot, M. Seymour, and R. Tipton. *Experimental characterization of a cold plate used in warm water cooling of data centers*. in *2017 33rd Thermal Measurement, Modeling & Management Symposium (SEMI-THERM)*. 2017.

78. Inc, C.S., *Operation and Maintenance Manual* 2015.

79. Sahini, M., U. Chowdhury, A. Siddarth, T. Pradip, D. Agonafer, R. Zeigham, J. Metcalf, and S. Branton. *Comparative study of high ambient inlet temperature effects on the performance of air vs. liquid cooled IT equipment*. in *Thermal and Thermomechanical Phenomena in Electronic Systems (ITherm), 2017 16th IEEE Intersociety Conference on*. 2017. IEEE.

80. Zimmermann, S., I. Meijer, M.K. Tiwari, S. Paredes, B. Michel, and D. Poulikakos, *Aquasar: A hot water cooled data center with direct energy reuse.* Energy, 2012. **43**(1): p. 237-245.

81. Zimmermann, S., 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, 2012. **55**(23): p. 6391-6399.

82. Addagatla, A., J. Fernandes, D. Mani, D. Agonafer, and V. Mulay. *Effect of warm water cooling for an isolated hybrid liquid cooled server*. in *Thermal Measurement, Modeling & Management Symposium (SEMI-THERM), 2015 31st*. 2015. IEEE.

83. Druzhinin, E., A. Shmelev, A. Moskovsky, Y. Migal, V. Mironov, and A. Semin. *High temperature coolant demonstrated for a computational cluster*. in *High Performance Computing & Simulation (HPCS), 2016 International Conference on*. 2016. IEEE.

84.   Druzhinin, E.A., A.B. Shmelev, A.A. Moskovsky, V.V. Mironov, and A. Semin, *Server Level Liquid Cooling: Do Higher System Temperatures Improve Energy Efficiency?* Supercomputing frontiers and innovations, 2016. **3**(1): p. 67-74.

85.   Garraghan, P., Y. Al-Anii, J. Summers, H. Thompson, N. Kapur, and K. Djemame. *A unified model for holistic power usage in cloud datacenter servers.* in *Proceedings of the 9th International Conference on Utility and Cloud Computing.* 2016. ACM.

86.   Zeighami, R., W.A. Saunders, H. Coles, and S. Branton. *Thermal performance modeling of hybrid liquid-air cooled servers.* in *Thermal and thermomechanical phenomena in electronic systems (itherm), 2014 ieee intersociety conference on.* 2014. IEEE.

87.   Li, L., W. Zheng, X. Wang, and X. Wang. *Placement optimization of liquid-cooled servers for power minimization in data centers.* in *Green Computing Conference (IGCC), 2014 International.* 2014. IEEE.

88.   Ovaska, S.J., R.E. Dragseth, and S.A. Hanssen. *Impact of retrofitted CPU water cooling on supercomputer performance and power consumption.* in *SoutheastCon, 2016.* 2016. IEEE.

89.   Ovaska, S.J., 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, 2016. **9**(3): p. 242-249.

90.   Kelkar, K.M. and S.V. Patankar. *Analysis and design of liquid-cooling systems using flow network modeling (FNM).* in *ASME 2003 International Electronic Packaging Technical Conference and Exhibition (InterPACK2003).* 2003.

91.   Ellsworth, M.J. *Flow network analysis of the IBM power 775 supercomputer water cooling system.* in *Thermal and Thermomechanical Phenomena in Electronic Systems (ITherm), 2014 IEEE Intersociety Conference on.* 2014. IEEE.

92.   Alkharabsheh, S., B. Ramakrishnan, and B. Sammakia. *Pressure drop analysis of direct liquid cooled (DLC) rack.* in *Thermal and Thermomechanical Phenomena in Electronic Systems (ITherm), 2017 16th IEEE Intersociety Conference on.* 2017. IEEE.

93. Sahini, M., C. Kshirsagar, M. Kumar, D. Agonafer, J. Fernandes, J. Na, V. Mulay, P. McGinn, and M. Soares. *Rack-level study of hybrid cooled servers using warm water cooling for distributed vs. centralized pumping systems.* in *2017 33rd Thermal Measurement, Modeling & Management Symposium (SEMI-THERM).* 2017.

94. Chu, R.C., M.K. Iyengar, V. Kamath, and R.R. Schmidt, *Energy efficient apparatus and method for cooling an electronics rack.* 2010, Google Patents.

95. Iyengar, M., M. David, P. Parida, V. Kamath, B. Kochuparambil, D. Graybill, M. Schultz, M. Gaynes, R. Simons, and R. Schmidt. *Server liquid cooling with chiller-less data center design to enable significant energy savings.* in *Semiconductor Thermal Measurement and Management Symposium (SEMI-THERM), 2012 28th Annual IEEE.* 2012. IEEE.

96. David, M.P., M.K. Iyengar, P. Parida, R.E. Simons, M. Schultz, M. Gaynes, R. Schmidt, and T. Chainer. *Impact of operating conditions on a chiller-less data center test facility with liquid cooled servers.* in *Thermal and Thermomechanical Phenomena in Electronic Systems (ITherm), 2012 13th IEEE Intersociety Conference on.* 2012. IEEE.

97. Iyengar, M., M. David, P. Parida, V. Kamath, B. Kochuparambil, D. Graybill, M. Schultz, M. Gaynes, R. Simons, and R. Schmidt. *Extreme energy efficiency using water cooled servers inside a chiller-less data center.* in *Thermal and Thermomechanical Phenomena in Electronic Systems (ITherm), 2012 13th IEEE Intersociety Conference on.* 2012. IEEE.

98. Parida, P.R., T. Chainer, M. Iyengar, M. David, M. Schultz, M. Gaynes, V. Kamath, B. Kochuparambil, R. Simons, and R. Schmidt. *System-Level Design for Liquid Cooled Chiller-Less Data Center.* in *ASME 2012 International Mechanical Engineering Congress and Exposition.* 2012. American Society of Mechanical Engineers.

99. Parida, P.R., T.J. Chainer, M.D. Schultz, and M.P. David. *Cooling Energy Reduction during dynamically controlled data center operation.* in *ASME 2013 International Technical Conference and Exhibition on Packaging and Integration of Electronic and Photonic Microsystems.* 2013. American Society of Mechanical Engineers.

100. Meyer, N., M. Ries, S. Solbrig, and T. Wettig. *iDataCool: HPC with hot-water cooling and energy reuse*. in *International Supercomputing Conference*. 2013. Springer.

101. Douchet, F., D. Nortershauser, S. Le Masson, and P. Glouannec, *Experimental and numerical study of water-cooled datacom equipment*. Applied Thermal Engineering, 2015. **84**: p. 350-359.

102. Schmidt, R., R.C. Chu, M. Ellsworth, M. Iyengar, D. Porter, V. Kamath, and B. Lehman. *Maintaining datacom rack inlet air temperatures with water cooled heat exchanger*. in *Asme 2005 pacific rim technical conference and exhibition on integration and packaging of mems, nems, and electronic systems collocated with the asme 2005 heat transfer summer conference*. 2005. American Society of Mechanical Engineers.

103. Schmidt, R., M. Iyengar, D. Porter, G. Weber, D. Graybill, and J. Steffes. *Open side car heat exchanger that removes entire server heat load without any added fan power*. in *Thermal and Thermomechanical Phenomena in Electronic Systems (ITherm), 2010 12th IEEE Intersociety Conference on*. 2010. IEEE.

104. Gao, T., E. Samadiani, B. Sammakia, and R. Schmidt. *Comparative thermal and energy analysis of a hybrid cooling data center with rear door heat exchangers*. in *ASME 2013 International Technical Conference and Exhibition on Packaging and Integration of Electronic and Photonic Microsystems*. 2013. American Society of Mechanical Engineers.

105. Wemhoff, A.P. and A. Ortega. *An exergy-based analysis of the effects of rear door heat exchange systems on data center energy efficiency*. in *Thermal and Thermomechanical Phenomena in Electronic Systems (ITherm), 2014 IEEE Intersociety Conference on*. 2014. IEEE.

106. Mulay, V., D. Agonafer, and R. Schmidt. *Liquid Cooling for Thermal Management of Data Centers*. in *ASME 2008 International Mechanical Engineering Congress and Exposition*. 2008. American Society of Mechanical Engineers.

107. Ellsworth, M., L. Campbell, R. Simons, M. Iyengar, R. Schmidt, and R. Chu. *The evolution of water cooling for IBM large server systems: Back to the future*. in *Thermal and Thermomechanical Phenomena in Electronic*

Systems, 2008. ITHERM 2008. 11th Intersociety Conference on. 2008. IEEE.

108. Schmidt, R. and M. Iyengar. *Server rack rear door heat exchanger and the new ASHRAE recommended environmental guidelines*. in *ASME 2009 InterPACK Conference collocated with the ASME 2009 Summer Heat Transfer Conference and the ASME 2009 3rd International Conference on Energy Sustainability*. 2009. American Society of Mechanical Engineers.

109. Almoli, A., A. Thompson, N. Kapur, J. Summers, H. Thompson, and G. Hannah, *Computational fluid dynamic investigation of liquid rack cooling in data centres.* Applied energy, 2012. **89**(1): p. 150-155.

110. Gao, T., M. David, J. Geer, R. Schmidt, and B. Sammakia, *Experimental and numerical dynamic investigation of an energy efficient liquid cooled chiller-less data center test facility.* Energy and Buildings, 2015. **91**: p. 83-96.

111. Sun Microsystems, I., *Sun Fire V20z and Sun Fire V40z Servers*. 2005.

112. Minyard, C., *IPMI–A Gentle Introduction with OpenIPMI*. 2004.

113. CoolIT. 2012; Available from: http://www.coolitsystems.com/.

114. Delta. *FFB 40 X 40 X 25 mm series*. 2017; Available from: http://www.deltaww.com/.

115. Al-Anii, Y.T.S., *Holistic and Energy-Efficient Management of Datacentres*. 2017, University of Leeds.

116. 3COM, *SuperStack® 3 Switch 3870 Family Getting Started Guide*, 3COM, Editor. 2004.

117. Airedale International Air Conditioning Limited. 2015; Available from: http://www.airedale.com/.

118. Wang, C.-C., J.-Y. Chang, and N.-F. Chiou, *Effects of waffle height on the air-side performance of wavy fin-and-tube heat exchangers.* Heat Transfer Engineering, 1999. **20**(3): p. 45-56.

119. Wang, C.-C. and K.-Y. Chi, *Heat transfer and friction characteristics of plain fin-and-tube heat exchangers, part I: new experimental data.* International Journal of Heat and Mass Transfer, 2000. **43**(15): p. 2681-2691.

120. Wongwises, S. and Y. Chokeman, *Effect of fin pitch and number of tube rows on the air side performance of herringbone wavy fin and tube heat*

exchangers. Energy Conversion and Management, 2005. **46**(13): p. 2216-2231.

121.   Youn, B. and N. Kim, *An experimental investigation on the airside performance of fin-and-tube heat exchangers having sinusoidal wave fins.* Heat and Mass Transfer, 2007. **43**(12): p. 1249-1262.

122.   Wang, C.-C., J.-S. Liaw, and B.-C. Yang, *Airside performance of herringbone wavy fin-and-tube heat exchangers–data with larger diameter tube.* International Journal of Heat and Mass Transfer, 2011. **54**(5): p. 1024-1029.

123.   Pongsoi, P., S. Pikulkajorn, C.-C. Wang, and S. Wongwises, *Effect of number of tube rows on the air-side performance of crimped spiral fin-and-tube heat exchanger with a multipass parallel and counter cross-flow configuration.* International Journal of Heat and Mass Transfer, 2012. **55**(4): p. 1403-1411.

124.   Yu, R., A. Sommers, and N. Okamoto, *Effect of a micro-grooved fin surface design on the air-side thermal-hydraulic performance of a plain fin-and-tube heat exchanger.* International Journal of Refrigeration, 2013. **36**(3): p. 1078-1089.

125.   Ma, X., G. Ding, Y. Zhang, and K. Wang, *Airside characteristics of heat, mass transfer and pressure drop for heat exchangers of tube-in hydrophilic coating wavy fin under dehumidifying conditions.* International Journal of Heat and Mass Transfer, 2009. **52**(19): p. 4358-4370.

126.   Min, J., X. Wu, L. Shen, and F. Gao, *Hydrophilic treatment and performance evaluation of copper finned tube evaporators.* Applied Thermal Engineering, 2011. **31**(14): p. 2936-2942.

127.   Wang, C.-C. and C.-T. Chang, *Heat and mass transfer for plate fin-and-tube heat exchangers, with and without hydrophilic coating.* International Journal of Heat and Mass Transfer, 1998. **41**(20): p. 3109-3120.

128.   Kim, H.-Y. and B. Kang, *Effects of hydrophilic surface treatment on evaporation heat transfer at the outside wall of horizontal tubes.* Applied thermal engineering, 2003. **23**(4): p. 449-458.

129.   Hwang, Y., S. Popli, and R. Radermacher, *Enhancement of Round Tube and Flat Tube-Louver Fin Heat Exchanger Performance Using Deluge*

Water Cooling, in *International Refrigeration and Air Conditioning Conference* 2012: Purdue University.

130. Bogan, N. and C. Park, *Effects of Tube Row on Heat Transfer and Surface Wetting of Microscale Porous-Layer Coated, Horizontal-Tube, Falling-Film Evaporator.* Journal of Heat Transfer (ASME), 2013.

131. Popli, S., Y. Hwang, and R. Radermacher. *Performance Enhancement of Enhanced Herringbone Wavy-Fin Round Tube Inclined Heat Exchangers With and Without Hydrophilic Coating Using Evaporative Spray and Deluge Cooling.* in *ASME 2013 International Mechanical Engineering Congress and Exposition.* 2013. American Society of Mechanical Engineers.

132. Wang, C., P. Chen, and J. Jang, *Heat transfer and friction characteristics of convex-louver fin-and-tube heat exchangers.* Experimental heat transfer, 1996. **9**(1): p. 61-78.

133. Wang, C., W. Fu, and C. Chang, *Heat transfer and friction characteristics of typical wavy fin-and-tube heat exchangers.* Experimental thermal and fluid science, 1997. **14**(2): p. 174-186.

134. Wang, C.-C., W.-S. Lee, W.-J. Sheu, and Y.-J. Chang, *A comparison of the airside performance of the fin-and-tube heat exchangers in wet conditions; with and without hydrophilic coating.* Applied Thermal Engineering, 2002. **22**(3): p. 267-278.

135. Liu, L., *Effects of air-side surface wettability on the performance of dehumidifying heat exchangers.* 2011, University of Illinois at Urbana-Champaign.

136. KG, e.-p.M.G.C.; Available from: www.ebmpapst.com.

137. ebmpapst, *EC axial fan - HyBlade.* 2015.

138. CAREL. *optiMist evaporative cooling* 2013; Available from: http://www.carel.com/.

139. GRUNDFOS. *CME 1-2 A-R-I-E-AQQE.* Available from: http://www.grundfos.com/.

140. Books, H., *Legionnaires' disease: Technical guidance.*

141. Clifton, J.R., W.J. Rossiter, and P.W. Brown, *Degraded aqueous glycol solutions: pH values and the effects of common ions on suppressing pH decreases.* Solar energy materials, 1985. **12**(1): p. 77-86.

142. Rossiter, W.J., M. Godette, P.W. Brown, and K.G. Galuk, *An investigation of the degradation of aqueous ethylene glycol and propylene glycol solutions using ion chromatography.* Solar Energy Materials, 1985. **11**(5-6): p. 455-467.

143. Note, T.A., *Traversing a Duct to Determine Average Air Velocity.* TI-106, 2014.

144. SPEC2008. *Standard Performance Evaluation Corporation.* 2017; Available from: https://www.spec.org/sfs2008/.

145. SPEC, S., *Power and Performance, Benchmark Methodology V2. 1.* 2011.

146. Voltech, *PM1000+ POWER ANALYZER USER MANUAL.* 2011.

147. Rossman, L.A., *EPANET 2: users manual.* 2000.

148. Cengel, Y.A. and J.M. Cimbala, *Fluid Mechanics. Vol. 1.* 2006, Tata McGraw-Hill Education, New York.

149. Xylem Inc. *Electronically commutated spherical motor with permanent magnet rotor/impeller.* Available from: www.xyleminc.com/brands/laingthermotech.

150. Gorti, N.P.K., *Application aware performance, power consumption, and reliability tradeoff.* 2014, Iowa State University.

151. Cader, T., V. Sorel, L. Westra, and A. Marquez, *Liquid Cooling in Data Centers.* ASHRAE Transactions, 2009. **115**(1).

152. Dhavaleswarapu, H.K., C.P. Migliaccio, S.V. Garimella, and J.Y. Murthy, *Experimental investigation of evaporation from low-contact-angle sessile droplets.* Langmuir, 2009. **26**(2): p. 880-888.

153. Kumari, N. and S.V. Garimella, *Characterization of the heat transfer accompanying electrowetting or gravity-induced droplet motion.* International Journal of Heat and Mass Transfer, 2011. **54**(17): p. 4037-4050.

154. Nguyen, T.A., A.V. Nguyen, M.A. Hampton, Z.P. Xu, L. Huang, and V. Rudolph, *Theoretical and experimental analysis of droplet evaporation on solid surfaces.* Chemical Engineering Science, 2012. **69**(1): p. 522-529.

155. Onan, C., D.B. Ozkan, and S. Erdem, *CFD and Experimental Analysis of a Falling Film outside Smooth and Helically Grooved Tubes.* Advances in Mechanical Engineering, 2014. **2014**.

156.   Wang, X. and A.M. Jacobi, *A Thermodynamic Basis for Predicting Falling-Film Mode Transitions.* International Journal of Refrigeration, 2014.

157.   Paterson, C., *Thin fluid films subject to external airflows.* 2013, University of Strathclyde.

158.   Wang, X., P. Hrnjak, S. Elbel, A. Jacobi, and M. He, *Heat transfer performance for a falling-film on horizontal flat tubes.* Journal of Heat Transfer, 2013. **135**(7): p. 072901.

159.   Jani, S. and M. Amini, *Heat Transfer Analysis of Falling Film Evaporation on a Horizontal Elliptical Tube.* Journal of Heat Transfer, 2012. **134**(6): p. 064505.

160.   David, S., K. Sefiane, and L. Tadrist, *Experimental investigation of the effect of thermal properties of the substrate in the wetting and evaporation of sessile drops.* Colloids and Surfaces A: Physicochemical and Engineering Aspects, 2007. **298**(1): p. 108-114.

161.   Mollaret, R., K. Sefiane, J. Christy, and D. Veyret, *Experimental and numerical investigation of the evaporation into air of a drop on a heated surface.* Chemical Engineering Research and Design, 2004. **82**(4): p. 471-480.

162.   Sobac, B. and D. Brutin, *Thermal effects of the substrate on water droplet evaporation.* Physical Review E, 2012. **86**(2): p. 021602.

163.   Shin, D.H., S.H. Lee, J.-Y. Jung, and J.Y. Yoo, *Evaporating characteristics of sessile droplet on hydrophobic and hydrophilic surfaces.* Microelectronic Engineering, 2009. **86**(4): p. 1350-1353.

164.   Gelderblom, H., A.G. Marin, H. Nair, A. van Houselt, L. Lefferts, J.H. Snoeijer, and D. Lohse, *How water droplets evaporate on a superhydrophobic substrate.* Physical Review E, 2011. **83**(2): p. 026306.

165.   Yu, Y.-S., Z. Wang, and Y.-P. Zhao, *Experimental and theoretical investigations of evaporation of sessile water droplet on hydrophobic surfaces.* Journal of Colloid and Interface Science, 2012. **365**(1): p. 254-259.

166.   Dash, S. and S.V. Garimella, *Droplet evaporation on heated hydrophobic and superhydrophobic surfaces.* Physical Review E, 2014. **89**(4): p. 042402.

167.   Sefiane, K., L. Tadrist, and M. Douglas, *Experimental study of evaporating water–ethanol mixture sessile drop: Influence of concentration.* International Journal of Heat and Mass Transfer, 2003. **46**(23): p. 4527-4534.

168.   Sefiane, K. and L. Tadrist, *Experimental investigation of the de-pinning phenomenon on rough surfaces of volatile drops.* International Communications in Heat and Mass Transfer, 2006. **33**(4): p. 482-490.

169.   Lee, K., C. Cheah, R. Copleston, V. Starov, and K. Sefiane, *Spreading and evaporation of sessile droplets: Universal behaviour in the case of complete wetting.* Colloids and Surfaces A: Physicochemical and Engineering Aspects, 2008. **323**(1): p. 63-72.

170.   Dunn, G., S. Wilson, B. Duffy, S. David, and K. Sefiane, *The strong influence of substrate conductivity on droplet evaporation.* Journal of Fluid Mechanics, 2009. **623**: p. 329-351.

171.   Stauber, J., S. Wilson, B. Duffy, and K. Sefiane, *On the lifetimes of evaporating droplets.* Journal of Fluid Mechanics, 2014. **744**: p. R2.

172.   Deegan, R.D., O. Bakajin, T.F. Dupont, G. Huber, S.R. Nagel, and T.A. Witten, *Contact line deposits in an evaporating drop.* Physical Review E, 2000. **62**(1): p. 756.

173.   Popov, Y.O., *Evaporative deposition patterns: spatial dimensions of the deposit.* Physical Review E, 2005. **71**(3): p. 036313.

174.   Dash, S. and S.V. Garimella, *Droplet evaporation dynamics on a superhydrophobic surface with negligible hysteresis.* Langmuir, 2013. **29**(34): p. 10785-10795.

175.   Gleason, K. and S.A. Putnam, *Microdroplet evaporation with a forced pinned contact line.* Langmuir, 2014. **30**(34): p. 10548-10555.

176.   Pan, Z., S. Dash, J.A. Weibel, and S.V. Garimella, *Assessment of water droplet evaporation mechanisms on hydrophobic and superhydrophobic substrates.* Langmuir, 2013. **29**(51): p. 15831-15841.

177.   Carle, F., B. Sobac, and D. Brutin, *Experimental evidence of the atmospheric convective transport contribution to sessile droplet evaporation.* Applied Physics Letters, 2013. **102**(6): p. 061603.

178.   Chen, Y., W. Hu, J. Wang, F. Hong, and P. Cheng, *Transient effects and mass convection in sessile droplet evaporation: The role of liquid and*

substrate thermophysical properties. International Journal of Heat and Mass Transfer, 2017. **108**: p. 2072-2087.

179.   Park, J.K., J. Ryu, B.C. Koo, S. Lee, and K.H. Kang, *How the change of contact angle occurs for an evaporating droplet: effect of impurity and attached water films.* Soft Matter, 2012. **8**(47): p. 11889-11896.

180.   Nguyen, T.A. and A.V. Nguyen, *Increased evaporation kinetics of sessile droplets by using nanoparticles.* Langmuir, 2012. **28**(49): p. 16725-16728.

181.   Pittoni, P.G., C.-C. Chang, T.-S. Yu, and S.-Y. Lin, *Evaporation of water drops on polymer surfaces: pinning, depinning and dynamics of the triple line.* Colloids and Surfaces A: Physicochemical and Engineering Aspects, 2013. **432**: p. 89-98.

182.   Putnam, S.A., A.M. Briones, L.W. Byrd, J.S. Ervin, M.S. Hanchak, A. White, and J.G. Jones, *Microdroplet evaporation on superheated surfaces.* International Journal of Heat and Mass Transfer, 2012. **55**(21-22): p. 5793-5807.

183.   Bormashenko, E., A. Musin, and M. Zinigrad, *Evaporation of droplets on strongly and weakly pinning surfaces and dynamics of the triple line.* Colloids and Surfaces A: Physicochemical and Engineering Aspects, 2011. **385**(1-3): p. 235-240.

184.   Blake, T. and J. De Coninck, *The influence of solid–liquid interactions on dynamic wetting.* Advances in colloid and interface science, 2002. **96**(1-3): p. 21-36.

185.   Anantharaju, N., M. Panchagnula, and S. Neti, *Evaporating drops on patterned surfaces: Transition from pinned to moving triple line.* Journal of colloid and interface science, 2009. **337**(1): p. 176-182.

186.   Kulinich, S. and M. Farzaneh, *Effect of contact angle hysteresis on water droplet evaporation from super-hydrophobic surfaces.* Applied Surface Science, 2009. **255**(7): p. 4056-4060.

187.   Stauber, J.M., S.K. Wilson, B.R. Duffy, and K. Sefiane, *On the lifetimes of evaporating droplets with related initial and receding contact angles.* Physics of Fluids, 2015. **27**(12): p. 122101.

188.   Hu, D. and H. Wu, *Volume evolution of small sessile droplets evaporating in stick-slip mode.* Physical Review E, 2016. **93**(4): p. 042805.

189. Sefiane, K. and R. Bennacer, *Nanofluids droplets evaporation kinetics and wetting dynamics on rough heated substrates.* Advances in Colloid and Interface Science, 2009. **147**: p. 263-271.

190. Bazargan, V., *Effect of substrate cooling and droplet shape and composition on the droplet evaporation and the deposition of particles*, in 2014, The University Of British Columbia. p. 120.

191. Srinivasan, S., G.H. McKinley, and R.E. Cohen, *Assessing the accuracy of contact angle measurements for sessile drops on liquid-repellent surfaces.* Langmuir, 2011. **27**(22): p. 13582-13589.

192. Lopes, M.C., E. Bonaccurso, T. Gambaryan-Roisman, and P. Stephan, *Influence of the substrate thermal properties on sessile droplet evaporation: Effect of transient heat transport.* Colloids and Surfaces A: Physicochemical and Engineering Aspects, 2013. **432**: p. 64-70.

193. Berthier, J., *Micro-drops and digital microfluidics.* 2012: William Andrew.

194. Oxford nanoSystems. Available from: http://oxfordnanosystems.com/.

195. *FTA4000.* Available from: http://www.firsttenangstroms.com.

196. Lebedev, N.N., *Special functions and their applications*, ed. R.A. Silverman. 1965, United States of America: Prentice-Hall, Inc. 322.

197. MATLAB, *Reference manual.* The Mathworks Inc. Natick, Massachusetts, United States, 2013.

198. Incropera, F.P., *Fundamentals of heat and mass transfer*. 2011: John Wiley & Sons.

199. Cengel, Y.A. and M.A. Boles, *Thermodynamics: an engineering approach.* Sea, 2002. **1000**: p. 8862.

200. Marrero, T.R. and E.A. Mason, *Gaseous diffusion coefficients.* Journal of Physical and Chemical Reference Data, 1972. **1**(1): p. 3-118.

201. Faghri, A. and Y. Zhang, *Transport phenomena in multiphase systems*. 2006: Academic Press.

202. Girard, F., M. Antoni, S. Faure, and A. Steinchen, *Influence of heating temperature and relative humidity in the evaporation of pinned droplets.* Colloids and Surfaces A: Physicochemical and Engineering Aspects, 2008. **323**(1): p. 36-49.

Appendix A - PID code to control the substrate temperature

```
/*********************************************************
 * PID code to control the substrate temperature
 *********************************************************/

#include <PID_v1.h>

//Define Variables we'll be connecting to
double Setpoint, Input, Output;

int sensorPin = 0; //the analog pin the TMP36's Vout (sense) pin is connected
to
                    //the resolution is 10 mV / degree centigrade with a
                    //500 mV offset to allow for negative temperatures
int heaterPin = 9;
int adjustedOutput; // to adjust the output of PID to arduino output (from 0
to 255)
//Specify the links and initial tuning parameters
PID myPID(&Input, &Output, &Setpoint,2,5,1, DIRECT);

void setup()
{

pinMode(heaterPin, OUTPUT);
  Serial.begin(9500);
  while (! Serial);
  Serial.println("Speed 0 to 255");
  //initialize the variables we're linked to
  //Input = temperatureC;
  Setpoint = 45;

  //turn the PID on
  myPID.SetMode(AUTOMATIC);
}

void loop()
{
      //getting the voltage reading from the temperature sensor
 int reading = analogRead(sensorPin);

 // converting that reading to voltage, for 3.3v arduino use 3.3
 float voltage = reading * 5.0;
 voltage /= 1024.0;

// now print out the temperature
 float temperatureC = (voltage - 1.25) / 0.005 ;  //converting from 10 mv per
degree wit 500 mV offset

                                    //to degrees ((voltage - 500mV)
times 100)
 Serial.print(temperatureC); Serial.println(" degrees C");

  Input = temperatureC;
  myPID.Compute();
  adjustedOutput = Output * 1.96;
  analogWrite(heaterPin,adjustedOutput);
}
```