DYNACOOL - SIMULATING EFFICIENT LIQUID COOLING FOR CURRENT AND NEXT GENERATION LARGE SCALE DATA CENTRES

A THESIS SUBMITTED TO AUCKLAND UNIVERSITY OF TECHNOLOGY
IN PARTIAL FULFILMENT OF THE REQUIREMENTS FOR THE DEGREE OF
MASTER OF COMPUTER AND INFORMATION SCIENCES

Supervisors

Dr. Roopak Sinha

06 March 2017

By

Nidhi Gowdra

School of Engineering, Computer and Mathematical Sciences

Copyright

Copyright in text of this thesis rests with the Author. Copies (by any process) either in full, or of extracts, may be made **only** in accordance with instructions given by the Author and lodged in the library, Auckland University of Technology. Details may be obtained from the Librarian. This page must form part of any such copies made. Further copies (by any process) of copies made in accordance with such instructions may not be made without the permission (in writing) of the Author.

The ownership of any intellectual property rights which may be described in this thesis is vested in the Auckland University of Technology, subject to any prior agreement to the contrary, and may not be made available for use by third parties without the written permission of the University, which will prescribe the terms and conditions of any such agreement.

Further information on the conditions under which disclosures and exploitation may take place is available from the Librarian.

© Copyright 2017. Nidhi Gowdra

Declaration

I hereby declare that this submission is my own work and that, to the best of my knowledge and belief, it contains no material previously published or written by another person nor material which to a substantial extent has been accepted for the qualification of any other degree or diploma of a university or other institution of higher learning.

Signature of candidate

Acknowledgements

I would first like to thank Dr. Roopak Sinha for being my thesis supervisor. His door was always open whenever I needed help or guidance about my research or writing. He consistently allowed this thesis to be my own work, but guided me patiently and calmly in the right direction whenever he thought I needed it. It's been a pleasure working with you and getting to learn from you on this thesis and on our additional projects.

Finally, I must express my very profound gratitude to my parents who have nurtured me and endured my tantrums all these years. I must thank my sister-in-law who has tolerated my antics and brother who has always looked out for me and stood by my side.

Thank you all for supporting this accomplishment, without whom this journey would not have been possible. Both personally and professionally this part of my journey has been the most challenging by far, but the most rewarding.

Looking forward to the next leg of my journey :D

Nidhi Gowdra, March 06th, 2017

Abstract

Energy consumption in Large Scale Data Centres (LSDC's) doubled from 2000 to 2006 reaching 61 TerraWatt-hour (TWh) per year. Most power generation sources are sadly fossil fuelled, which is increasing the effects of anthropomorphic climate change. 99-100% of the energy consumed by IT equipment is dissipated as heat from the servers, which creates a real problem of cooling in LSDC's. Air cooling systems in LSDC's are struggling to handle the increased cooling demands, which is why they account for 40% of the total energy consumed. The rest of the energy consumed is used to power the IT equipment and data centre infrastructure facilities like lighting. The ratio of power consumed by the data centre facility to the power consumed by the IT equipment is known as Power Usage Effectiveness (PUE), which is a metric used to measure data centre efficiency. Reducing the PUE by even fractions of percentages can prevent millions of tons of greenhouse gases from being emitted into the atmosphere.

One of the methods of reducing PUE is by using alternate forms of cooling technologies like liquid cooling. This thesis explores novel optimisation methods for cooling control in liquid cooled LSDC's. The three strategies focused on are Static Flow Rate, Variable Flow Rate (VFR) and the proposed Pulsed Variable Flow Rate (PVFR) cooling control strategies. The power consumption of coolant pumps and the effect it has on reducing energy consumption and PUE are investigated for all three cooling control strategies using computer simulations. Current simulation software were limited to air cooling and as such we needed to develop a proprietary computer simulation software.

The software we developed was *DynaCool* and the simulation data we gathered was used to analyse the effectiveness of the different cooling control strategies.

DynaCool was built using requirements engineering and model driven design to ensure the validity of the software, as these methodologies are commonly used in large complex industrial systems. The data analysed indicated a PUE reduction of at least 15.4% for the novel PVFR liquid cooling control strategy over the static and VFR cooling control strategies. This reduction equates to savings of 2.84 million tons of greenhouse gas emissions and 18.788 TWh of power consumption per year for an adoption rate of 100%. Realistically speaking, an adoption rate of 10% would yield power saving of 1.88 TWh or 284 thousand tons of greenhouse gas emission per year. This adoption rate is easily achievable by the industry as recent trends indicate data centre operators are moving towards alternate cooling technologies. This is evidenced by Google aiming to achieve carbon neutrality by 2017 using liquid cooling technologies.

Contents

List of Tables

List of Figures

Chapter 1

Introduction

1.1 Motivation

The rapid adoption of electronics in all sectors of the world has made it one of the most influential industries in history. A few market leaders like Intel, Microsoft and Apple among others are driving innovation (Gawer & Cusumano, 2002). To name just one, cell phones are the fastest and largest adoption of a technology from introduction to mainstream use in human history (Jannou, 2015). These innovations require computing infrastructure to host services and data, which are provided by data centres which consume large amounts of power. It was approximately 61 TerraWatt-Hour (TWH) in 2006 for the USA alone (R. Brown et al., 2008). Which equated to 1.5 % of the total energy use of the country. The main sources of this power are fossil fuels, which is leading to an increasing amount of green house gas emissions per year (Oreskes, 2004).

Global climate change is *the most crucial issue* facing our generation and the next. Anthropomorphic climate change is melting polar ice caps at an alarming rate, causing a mass extinction of species (Stocker, 2014). Due to such dire consequences, there have been urgent calls made to reduce greenhouse gas emissions in all industries (Sultan, 2010). The high-technology industry has been heavily targeted, and is trying to mitigate it's impact by reducing emissions (Conservative & coalition government of the UK, 2016). Data centre operators' prime objective is to reduce power consumption as this would yield the greatest reductions in carbon emissions other than increasing Information Technology (IT) equipment life cycles (Webb, 2008). Data centre operators are increasingly looking to reduce power consumptions of cooling systems by increasing their operating efficiencies, this is because the cooling systems account for 40 % of the total power consumed in data centres (Capozzoli & Primiceri, 2015), while the rest is used to power the facility and IT equipment.

1.2 Domain of research

Data Centres are massive facilities hosting a large number of servers (hundreds to tens of thousands) and supporting infrastructures needed for continued operation. Large scale data centres (LSDC's), according to Patel, Sharma, Bash and Beitelmal (2002) consist of computing components and infrastructure components which contribute to the thermal load. The servers are arranged in server racks which have been standardised to allow for interoperability amongst the components. The density of server racks are measured in KiloWatts (KW) per server rack which indicate the power consumption (Hastings, Varghese, Jasso & Leigh, 2002).

The computing components in LSDC's are computer servers, storage devices like hard drives, solid state drives and computer memory, networking devices like switches and routers and miscellaneous on-board components (Sanders, legal representative Vivian, Nguyen, Pascarella et al., 2006). The infrastructure components are Power Distribution Units (PDU's), Uninterruptible Power Supplies (UPS's), Computer Room Air-Conditioners (CRAC's), fans, chillers and cooling towers (Patterson & Fenwick, 2008; Patel et al., 2002).

The demand for cloud computing resources has been largely driven by consumers.

Recently however, public cloud enterprise information systems have begun gaining widespread adoption due to their financial and functional benefits, as noted by Chang and Hsu (2016). With the expeditious adoption of cloud technologies, there is an urgent need for more computing resources in data centres, leading data centre operators to increase their server rack densities. Server rack power densities have increased from under 1KW per server rack to 10KW or more per server rack (Patel et al., 2002). Ten years later when Musilli and Ellison (2012) investigated the power densities of current modern LSDC's, they found that future LSDC's are being designed to operate at 100 KW per server rack or more. This data indicates that LSDC's are holding true to the predicted trends and are consuming power at unprecedented rates, in the USA alone the data centre industry consumed 61 billion KWH or 61 TWH of electricity in 2006 (R. Brown et al., 2008).

According to Shi, Wan, Yan and Suo (2011), Cyber Physical Systems (CPS's) are characterised by computational and physical processes interacting in large complex systems. CPS's have sensors and feedback loops to allow for a cause and effect relationship between the computational and physical processes. A data centre satisfies all the requirements to be classified a CPS. Servers are affected by the computational load placed on them, which in turn effects the thermal profile of the data centre. This interaction is just one among many, and could be affected by other processes, which in turn may effect, or be affected by other such processes. Cooling control systems with thermal management systems are used to optimise the operating parameters of the servers.

1.3 The problem

The challenges facing data centres are many and are discussed in the sections below (refer Sections 1.3.1, 1.3.2, 1.3.3, 1.3.4 and 1.3.5). We only focus on the thermodynamic aspect of the data centres, this is because optimising cooling systems in LSDC's prove to offer the greatest opportunities for contributions (refer Section 1.5).

1.3.1 The cooling challenge

Increased server rack power densities create a real problem for data centre operators in terms of cooling. This has been the case ever since the days of ENIAC, the first electronic computer. ENIAC was cooled using air (Simons, 1995), but the problem with traditional air cooling is that it's highly inefficient. Liquid cooling technologies offer better efficiencies when compared to air cooling (Dai, Ohadi, Das & Pecht, 2014).

Cooling had not been considered an issue as early computers weighed tons, and occupied the size of large rooms whilst costing millions of dollars. This, dwarfs the cost of cooling significantly (McCartney, 1999) and as such was neglected. This all changed with a new innovation called the transistor, which led to a large scale adoption of computing and ushered in the 'digital revolution'. This is because transistors used a fraction of the power of large inefficient vacuum tubes, which soon made them obsolete (Koomey, Berard, Sanchez & Wong, 2011).

Due to the benefits offered by the transistors, computers increased in power densities significantly, as large amounts of these small transistors could be placed in a small space. The increase in density and the demand from consumer electronics during the digital revolution facilitated the introduction of mainframe computers. Mainframe computers are high density computers used to meet large scale demands for computing resources, their reminiscent technologies are still being used in data centre servers even today (Eustis, 2009).

The problem with mainframes however, was that they had high thermal outputs and air cooling was not sufficient to cool such high thermal loads (Schmidt, 2004). This is why they were shipped with liquid cooling technologies to accommodate the increase in thermal loads (Simons, 1995). But mainframes became obsolete with the introduction of microprocessors by Intel (Gawer & Cusumano, 2002), and these first microprocessors offered significantly lower power densities, thereby negating the challenge of high density cooling. The problem of cooling has returned with the unprecedented increase in server rack power densities and power consumptions (R. Schmidt, 2005), which use a large collection of these microprocessors. Currently, cooling systems in data centres account for approximately 40% of total power consumed (Capozzoli & Primiceri, 2015). This high power consumption has led data centre operators to enhance the efficiency of their cooling systems (Capozzoli & Primiceri, 2015; Jinkyun, Lim & Kim, 2009) or adopt alternate cooling technologies to reduce power consumptions (Facebook, 2010; Google, 2016a; 3M, 2016).

Metrics used to measure efficiency

The metrics used to measure and compare power consumptions of data centres are Power Usage Effectiveness (PUE) and Data Centre infrastructure Efficiency (DCiE) (refer section 2.1.1). PUE is the ratio of the power consumed by the data centre facility to the power consumed by the IT equipment and DCiE is the reciprocal of PUE, which is expressed as a percentage. Generally speaking, a PUE or DCiE of 3.0 & 33% respectively is highly inefficient and a PUE or DCiE of 1.2 & 83% respectively is highly efficient as ascertained by Belady, Rawson, Pfleuger and Cader (2008). A PUE of 2.0 implies, 2 KW of power is consumed by the data centre facility to power 1 KW of IT equipment.

Impact

Capozzoli and Primiceri (2015) state that modern LSDC's are at a tipping point when it comes to the cooling capabilities of air. This statement is seen as an opinion, because there is no consensus which states the limit of air cooling effectiveness. That being said however, alternative solutions are being preferred as they can cool larger heat loads more efficiently (Askwig, Butler, Huglen & Sinda, n.d.). Many different forms of cooling systems exist, each having their own efficiencies, and the metrics used to measure and compare power efficiencies have been explained previously (refer section 1.3.1).

To enhance cooling efficiencies, current and future LSDC's are moving to liquid cooled or hybrid cooled technologies. With several of the tech giants like Google already adopting liquid cooling in their data centres (Google, 2017b, 2016b). Other LSDC operators like Facebook have opted to implement a hybrid cooling approach (Facebook, 2010), which utilises both air and liquid cooling technologies. These solutions are explained in later sections of this thesis (refer Section 1.3.2).

We have chosen liquid cooled data centres as a key area for improvement because it is one of the only methods of cooling that offers a PUE of 1.2 or less, meaning it is highly efficient (3M, 2016; Google, 2017a). Especially when compared to air cooling which offers at the very best, a PUE of 1.4 (Dai et al., 2014). Liquid cooling in LSDC's is an emerging technology, and as such the systems are still in their first or second generations. This provides an opportunity to offer improvements (Dai et al., 2014) in liquid cooling technologies. Air cooling technologies on the other hand have already matured and have been optimized for peak performance.

Reducing PUE by even fractions of percentages will equate to large energy savings which reduces the carbon footprints of data centres, which saves tons of greenhouse gas emissions each year. This is because LSDC's operate at enormous scales, making savings of even fractions equate to large overall savings. This aim of reducing power consumption in LSDC's is in accordance with our overarching goal to reduce the impact of anthropomorphic climate change.

Size of the problem?

According to Felter et al. (2003), increased power densities offer greater performance, but at the cost of unprecedented power consumption and heat loads at both rack, and facility levels. Furthermore, Capozzoli and Primiceri (2015) explain that the main reason for high thermal loads is the fact that IT equipment converts around 99 - 100% of the power they consume into heat. This means that LSDC's with thousands of server racks each consuming 100 KW per server racks will expel enormous amounts of heat which needs to be dissipated efficiently (Musilli & Ellison, 2012).

Koomey (2011) asserts that, the power consumptions of data centres have roughly doubled between 2000 to 2005, and have continued to increase by 56% between 2005 to 2010. This rapid increase in power consumption has placed heavy stresses on cooling systems which are struggling to meet the cooling demand placed by large heat loads. The problem of cooling is exacerbated when LSDC's employ hundreds or thousands of these high density server racks and place them in a small space to make operating LSDC's financially feasible (Jinkyun et al., 2009).

1.3.2 Existing solutions

There have been many implementations of different cooling systems, and they are classified based on the active medium (fluid) used for cooling. The different cooling technologies are air, liquid and hybrid cooling systems.

Air cooling

Air cooling has been used ever since the inception of modern computing in 1946, when the first computer, the ENIAC was introduced (McCartney, 1999). One of the main advantages of air cooling is its simplicity of implementation, especially in smaller scale data centres. Dai et al. (2014) states that using air cooling is justified if the server racks have low power densities as it would be advantageous to implement air cooling over other forms of cooling.

Liquid cooling

Dai et al. (2014) asserts that liquid cooling is an emerging technology and as such it has seen very sparse adoption in the data centre industry. Google being a predominant player in the data centre industry, has already adopted liquid cooling in all of their data centres (Google, 2016a). Google's aim is to use sustainable approaches for cooling their fleet of data centres and are on the verge of achieving carbon neutral operation by 2017 (Google, 2016b). As alluded to earlier, this shift to sustainable cooling is in part due to increased calls to reduce greenhouse gas emissions by governments (Conservative & coalition government of the UK, 2016). Other corporations are investigating the merits of liquid cooling technologies but have been slow to adopt them (3M, 2016; Madhavi, 2014).

Hybrid cooling

Hybrid cooling coalesces air and water cooling technologies and is one of the methods of cooling LSDC's proposed by Lin and Ponnappan (2003). The technology utilises a fine mist of water sprayed into freely moving air, thereby leveraging the benefits of both technologies for enhanced cooling. Facebook has already implemented such a hybrid approach to cool their data centres, but other researchers and corporations remain sceptical of such a technology for cooling electronic equipment.

Prevalence of air cooling

As stated earlier, air cooling has been used for a long time and as such air cooling technologies have had time to mature and increase their efficiency by utilising air-side economisers to reduce power consumption (Dai et al., 2014). Air cooling technologies have been optimised for peak performance and are especially feasible for smaller scale data centres using lower power densities in their server racks. Thus due to the simplistic nature and prevalence of legacy data designs, air cooling has been dominant in adoption and utilisation in most data centres. As server racks power densities continue to increase at an accelerated pace, the power consumed by these systems are leading data centre operators to adopt alternate cooling technologies (Askwig et al., n.d.).

1.3.3 Limitations of cooling technologies

Air cooling

Air cooling uses a lot of energy, this is because air, as a cooling medium is less efficient when compared to using other forms of cooling medium like liquids in high heat load applications. Water as a cooling medium is 3467 times more efficient at removing heat than air (Aquatherm, 2015).

The bottleneck of air cooling technology is that, when power densities increase, the energy consumed for cooling increases exponentially (Shrivastava, Sammakia, Schmidt, Iyengar & Van Gilder, 2006). The causes for this are discussed at length in later chapters of this thesis (refer section 2.2.4). When Ebrahimi, Jones and Fleischer (2014) reviewed air cooling with alternate cooling technologies, they found that air cooling systems are more expensive to operate than liquid cooling systems. They also found that while the majority of legacy data centres still use air cooling systems, many LSDC operators are adopting alternate cooling technologies in current and future data centres. This is also evidenced by the implementations made by Facebook and Google (Facebook, 2010; Google, 2016a).

Recently, 3M (2016) achieved a PUE of 1.01 in testing, using a liquid immersion technology, one of the many forms of liquid cooling technologies. This makes such an implementation of liquid cooling technologies the most efficient form of cooling in LSDC's. This result is for a small subset of an actual LSDC and thus should not be extrapolated to be indicative of fleet-wide PUE's. Air cooling achieves at best a PUE of 1.4 whilst using expensive air-side economisers (Dai et al., 2014).

Hybrid cooling

While this approach is adequate to cool LSDC's efficiently, they require complex climate and thermal management systems to ensure safe operation of electronic components (Clark, 2013). Klein, Singh, Schappert, Griffel and Hamann (2011) have investigated the negative consequences of using such systems in data centres. They found that incorrect humidity management leads to corrosion and decrease the lifespans of electronic systems. This is why strict guidelines for data centre operators have been published by the American Society of Heating, Refrigeration and Air-Conditioning Engineers (ASHRAE) (2011).

1.3.4 Liquid cooling control strategies

Google's innovative technologies like using ocean water to help facilitate liquid cooling (Google, 2017b), along with complex Artificial Intelligence (AI) systems help reduce their cooling costs by up to 40% (Google, 2017a) or PUE by up to 15%. This adoption in liquid cooling causes them to have the lowest fleet-wide PUE of 1.11 (Google, 2016a). This is achieved by using dynamic cooling control strategies to vary the mass flow rate of the coolant according to the demands placed on servers. Mass flow rate is the amount (mass) of coolant flowing through the pipes per second of operation, expressed as kg/s.

There are two main cooling control strategies, which either vary the coolant continuously, known as a Variable Flow Rate (VFR) cooling control strategy or use a constant coolant flow, known as Static Flow Rate cooling control strategy. These control strategies are discussed in later chapters of this thesis (refer Section 2.5).

Although AI systems offer a significant reduction in power consumption and PUE's, they have a key limitation. According to Gao (2014), AI systems rely heavily on the accuracy of sensors and sensor data, which is essential for effective operation. Another limitation identified is that these systems require large amounts of computing resources which adds to the operating overhead costs. Cooling control strategies have a significant impact on reducing PUE's, and as such we aim to address the problems listed above, whilst optimising the cooling control strategies.

1.3.5 Data centre topologies

Topology refers to the interconnection or arrangement of the various different components of a system. In liquid cooled LSDC's, they represent the fluid flow dynamics between the various different components (refer Section 2.4.1). The topological configuration of the various components in LSDC's have a considerable impact on the heat profile on the data centre. This is because of the complex molecular interactions that occur in thermodynamics, which is explained in later chapters of this thesis (refer Section 2.2.4).

The most commonly used topology for interconnecting servers and server racks in liquid cooled systems is daisy chaining. Daisy chaining is where the coolant input to a server is the coolant output from the previous server (Madhavi, 2014). This topological configuration although convenient and offers lower installation and maintenance costs, it has a negative effect on efficiency. Daisy chaining reduces the efficiency of liquid cooling significantly when compared to other configurations, which is explained in later chapters of this thesis (refer Section 2.4.1). This is why liquid cooling technologies are particularly sensitive to increases in the number of server racks.

The other parameter that impacts the efficiency of all cooling technologies is the server rack power density. Increases in the heat load will have an increase in cooling demand. Liquid cooling technologies however, are particularly robust in accommodating the increasing changes in server rack power densities. This is because of the thermodynamics involved in liquid cooling, where the problem of efficiency is not predominantly affected by the increase in power densities but rather by increases in the number of server racks.

1.4 Research Questions

1.4.1 Research Question 1 (RQ 1)

'What are the key functional and non-functional requirements for cooling systems in next generation LSDC's?'

We opt for a Systematic Literature Review (SLR) to search existing literatures in order to determine the cooling requirements for next generation LSDC's. The process is detailed in later chapters of this thesis (see Section 2.1). We have opted for a SLR as it minimises biases and allows for a systematic search of a large research area.

1.4.2 Research Question 2 (RQ 2)

'How can we design more efficient cooling systems for next-generation LSDC's such that the key functional and non-functional requirements are met?'

After analysing the gaps and determining the limitations of current cooling technologies, we can propose a novel method for optimising liquid cooling technologies. We aim to model exiting solutions and test these with our proposed solution using Model Driven Engineering (MDE) (refer Section 4.2). According to Smiałek and Nowakowski (2015), complex systems like the thermal management systems used in LSDC's are built using a MDE approach. This is because MDE allows for abstraction of processes into manageable sub-units. Using a MDE approach also allows for the same implementation of control strategies across different platforms, which is the case for data centres operators as they are optimised for a particular operating location (Jinkyun, Lim & Kim, 2012).

1.5 Proposing a solution

In the literature review chapter of this thesis (refer Section 2.6), we identify an opportunity to optimise liquid cooling systems by analysing the gap in knowledge. To summarise, the gap in cooling systems is that they are either statically cooled, as is the case proposed by (3M, 2016), or employ a continuously varying coolant flow, as is the case of Google (Google, 2017a). We hypothesise that using pulses of intermediate coolant flow rather than employing continuously varying or static coolant flow will reduce power consumption of the cooling systems whilst maintaining safe operating parameters. The solution we propose is known as Pulsed Variable Flow Rate (PVFR) cooling control strategy. This solution is achieved using pulse width modulation (PWM), explained in later chapters of this thesis (refer Section 2.6.3). The aim of PVFR is to reduce the power consumption of coolant pumps by operating the pumps at their maximum efficiencies and exploiting the increased resistance of thermal fluctuations by modern electrical components (Intel, 1998).

1.5.1 Hypothesis testing

To test our hypothesis, we will use simulation software as it provides various benefits over physical implementations (refer Section 2.6.4). The requirements were analysed using a Requirements Engineering (RE) process explained in later chapters of this thesis (refer Section 4.1). As there is no commercially available simulation software which satisfy our requirements it led us to develop 'DynaCool', the first liquid cooling simulation software. To ensure the results are devoid of any biases we use statistical data analysis methods such as t-tests to analyse the data gathered from DynaCool (refer Section 6.3).

1.6 Contributions

By completing this research, we aim to deliver the following value additions to the research domain and consider this research a success if we are able to achieve all of the following contributions.

1. DynaCool (refer chapter 5): To develop an accurate simulation software for analysing liquid cooling in data centre server racks using computational fluid dynamics (CFD) calculations. DynaCool needs to be a modular system such that different modules can handle different loads and adjust flow rates and other characteristics based on peak demands. This is based on the data centre designs proposed by Geet (2014).

2. PVFR cooling control strategy: To develop an effective cooling control strategy using proprietary PVFR algorithm for dynamic liquid cooling.

3. Analysis of PVFR against VFR cooling control strategies (refer chapter 6): To achieve a statistically significant power reduction in coolant pumps used for liquid cooling.

1.6.1 DynaCool

We aim to develop DynaCool, a proof of concept application created to test our PVFR hypothesis, it needs to be a liquid cooling simulation engine for a server rack level of simulation accuracy. DynaCool is developed out of necessity, as there is no commercial software which meet the requirements of liquid cooling in LSDC's. Validation and accuracy of the software has to be thoroughly tested and verified, which is explained in later chapters of this thesis (refer Section 5.2.3). Iterative development is commonly used in large complex software development projects to reduce cost and complexity (Wolverton, 1974). Companies usually ship minimum viable products and then add more features in subsequent iterations based on various criteria (Moogk, 2012). This is similar to our approach of using versions for the DynaCool software.

1.6.2 Simulation models

We intend to provide a simulation model for further use by researchers looking to test, verify and validate our findings. These models help future researchers conduct better analyses into further optimisation in liquid cooling LSDC. The link to download the raw data along with the DynaCool software is available in the Appendix A.1.

1.6.3 Optimised liquid cooling control strategies

The main goal of this thesis is to propose and develop a novel PVFR liquid cooling algorithm and control strategy aimed at data centre analysts with the intention that they will be implemented in physical data centres. While the proposed PVFR algorithm is initially calibrated for liquid cooling systems in LSDC's they could be applied to any industry where mission critical cooling of high heat load equipment is required.

1.7 Thesis layout

Chapter 2, the literature review is done to identify gaps and opportunities in the domain of research. The research domain addresses the cooling challenge in LSDC's. Chapter 3, Research Methodology outlines the approaches we take whilst undertaking the research and developing solutions to address the identified problems. Chapter 4, focuses on system design using requirements modelling identified through requirements engineering. Chapter 5 outlines the process of building DynaCool using a MDE approach. Chapter 6 discusses the experimental design and the application of valid data analysis methods to identify the effectiveness of our PVFR hypothesis against existing cooling control strategies. Finally, Chapter 7 provides a summary of the research, possible directions and avenues for future research and concludes this thesis.

Chapter 2

Literature Review

The purpose of this chapter is to critically review existing literature for our chosen research domain i.e. cooling in LSDC's. We perform an exhaustive literature review with the intention of identifying any gaps in the current understanding of cooling in LSDC's. These gaps provide an opportunity to make contributions for advancing the field of cooling in LSDC's.

Section 2.1 illustrates the systematic literature review process used, while section 2.2 discusses the challenges faced by LSDC operators. Section 2.2 also discusses the cooling challenges and the relationship of proper cooling with some of the other challenges faced by LSDC operators. Section 2.3 focuses on exploring the methods of data centre cooling and their limitations. Section 2.4 examines the factors involved in the selection of cooling systems in LSDC's. Section 2.5 analyses the cooling control strategies employed in current generation LSDC's and their limitations. Section 2.6 highlights the findings from the exhaustive literature search to address the research questions and identify any gaps in knowledge to illustrate any research opportunities. Section 2.7 provides a conclusion for the chapter.

2.1 Systematic Literature Review (SLR) process

A well written literature review serves to provide an overview of the concepts, ideas and technologies currently being used. Webster and Watson (2002) asserts that an effective literature review is one which creates a strong foundation for advancing knowledge in the chosen research domain. Cronin, Ryan and Coughlan (2008) states that a good literature review is one which removes personal biases while gathering information from many credible sources. Whilst this statement is true, biases are often introduced unknowingly. Therefore, to minimise any biases we have chosen to undertake a systematic literature review (SLR) which offers several benefits over other literature review methods.

According to Brereton, Kitchenham, Budgen, Turner and Khalil (2007) one of the main benefits of SLR's over other methods of literature reviews is that, SLR's eliminate the inevitability of limited scope which ensue in other review methods. SLR's offer a systematic way to locate and assess relevant research in the chosen domain. It also offers to narrow down and segregate complex, vast research topics into manageable units. Cronin et al. (2008) suggests that SLR's should be used in applications where critically evaluating relevant literature is a priority, which is the case in our thesis.

According to B. A. Kitchenham, Dyba and Jorgensen (2004) for a software engineering discipline the process of performing an SLR is as follows,

1. Systematic literature review process and
2. Literature results

2.1.1 Systematic literature review process

The literature review process is conducted based on the guidelines outlined by B. Kitchenham et al. (2009), which are as follows.

1. Identify research questions

2. Search literature in the relevant domain
3. Apply inclusion and exclusion criteria to narrow down the peer reviewed papers applicable to our research area
4. Perform a quality assessment on the subset of papers based on certain Quality Assessment (QA) criteria
5. Data collection
6. Data analysis

Research questions

In order to critically develop the research questions, it is first necessary to explain the purpose, focus and problem definition of the research being undertaken.

Purpose of Research

The demand for cloud computing has been growing rapidly and data centre operators are struggling to keep up with the demand (Barroso, Clidaras & Hölzle, 2013). Data centres consumed 61 TWH of electricity in the US alone (Capozzoli & Primiceri, 2015). The net effect of our contribution by reducing even fractions of percentages has the impact to save millions of dollars in direct costs from reduced power consumption. The hope of this research however, is that if implemented, it will save millions of tons of greenhouse gases from being emitted into the atmosphere thereby saving billions of dollars in indirect costs for carbon mediation. *This serves as both an environmental and economic incentive to adopt the solutions outlined in this thesis.* The purpose of this research is to reduce energy consumption of cooling pumps by optimising the cooling control strategies used in liquid cooled LSDC's.

Metrics used to measure efficiencies in LSDC's

According to Belady et al. (2008), Power Usage Effectiveness (PUE) is a metric that is widely used by the data centre industry and is ubiquitous in its adoption to measure power consumption efficiencies (Jinkyun et al., 2012). PUE is a ratio of the total power consumed by the data centre facility to the power consumed by the IT equipment.
i.e. PUE = Total facility power / IT equipment power

Total facility power includes all of the supporting infrastructures such as the lighting, power distribution, backup power, cooling systems etc. which are critical for operation. IT equipment power refers only to the power consumed by the servers, network switches, and other computer components such as displays, monitoring equipment etc. Data Centre Infrastructure Efficiency (DCiE) is another metric that is commonly used but not widely adopted to measure efficiencies in data centres. DCiE is the reciprocal of PUE and is expressed in percentages.
i.e. DCiE = 1 / PUE = IT equipment power / Total facility power × 100 (%)

Both metrics illustrate the same measure of power consumption efficiencies. For example, a PUE of 3.0 (DCiE of 33%) indicates the data centre facility consumes three times as much power as the IT equipment.

2.1.2 Research focus

We recognise the significance of optimising cooling control strategies for reducing power consumptions of cooling pumps to increase the PUE of LSDC's. This will reduce power consumption of data centres and thus minimise the impact of anthropomorphic climate change by reducing greenhouse gas emissions of data centres.

We propose to optimise liquid cooling control strategies by implementing control logic on Programmable Logic Controller's (PLC's). PLC's are the basic control mechanisms in large cyber physical systems such as data centres, and they use significantly less computing resources to implement when compared to other control systems, which is why they are heavily used. PLC's manage a cyber physical system's infrastructures and control systems, which is why the malware STUXNET had such a devastating impact on Iranian nuclear enrichment centrifuges (Karnouskos, 2011).

Problem definition

Power densities of server racks have been increasing rapidly, which was evidenced when Madhavi (2014) investigated the performance of different cooling systems in LSDC's. According to R. Schmidt et al. (2005), legacy data centres operate at power densities of 15 kW per server rack and modern data centres are currently operating at power densities of 30 kW per server rack. These power densities dwarf the scale of future LSDC's, which are being designed to operate at power densities of 100 kW per server rack (Musilli & Ellison, 2012).

Higher density server racks have increased power consumption requirements and as such they possess larger thermal dissipations. According to Capozzoli and Primiceri (2015), this is because 99-100% of the power consumed by servers is expelled as heat. To ensure safe operation the thermal dissipation needs to be managed carefully, which leads us to define our research problem,

'Considering the requirements of future next generation LSDC's, what would be an efficient way to cool the server racks?'.

Research questions can now be formulated as from the research problem which are as follows,

RQ 1. *'What are the key functional and non-functional requirements for cooling systems in next generation LSDC's?'*

RQ 2. *'How can we design more efficient cooling systems for next-generation LSDC's such that key functional and non-functional requirements are met?'*

Each of these research questions will be discussed comprehensively in subsequent chapters (refer Chapters 3, 4, 5, 6 and 7). We leverage the use of engineering principles and 'the scientific process' to reduce any personal biases to answer the research questions.

Research scope

This research is limited to evaluate only liquid cooling systems and their control strategies employed in LSDC's. This is because, the solutions and contributions outlined in this thesis will have the most impact in such operating environments (refer Section 1.5). Although this is the case, the contributions made in this thesis are designed to be effective on smaller scale and lower density data centres. The thesis, in its current form focuses mainly on thermodynamic properties using computer simulations. The models used in the simulations are modelled using computer science disciplines like requirements engineering and model driven design (refer Chapter 4).

2.1.3 Search process

According to Cronin et al. (2008), searching for relevant literature can be done either manually or using automated programs, the latter being used by the majority of researchers. While an automated search for literature using only keywords and titles is an efficient method of performing an SLR, it misses several important papers as evidenced by B. Kitchenham et al. (2009). They performed an automated search and compared the results with one done manually. It was found that when an automated search was done for a period from 2004 to 2007 and then for only the year 2008, using the same criteria for both manual and automated searches, the automated search missed 3 relevant papers of which they were aware.

Doing a manual search however, is time intensive and results in reviewing only a fraction of the literature available in the domain. This is why we have opted for a hybrid approach, utilising both automated and manual searches. This approach offers the time to review literatures from a wider spectrum in the research domain.

The search process we employed was done using an automated search of keywords using Google Scholar and specific highly rated databases such as IEEE Computer Society Digital Library, SpringerLink and ScienceDirect. These databases were specifically chosen because they are highly ranked by journal ranking portals such as CORE, an Australian based ranking service. We also used Google web search to search for relevant whitepapers using specific keywords.

The keywords were based on analysing the words occurring in titles and abstracts along with a few wild card keywords to cover a broad search domain. Google web search was also used to search for relevant white papers and other secondary sources of paper by using data centre equipment manufacturers websites and/or their products as the keyword. For the full list of the keywords used, please refer Appendix A.2.

The manual search was done based on reviewing specific conferences and journals such as journal of systems and software, IEEE transactions on software engineering, IEEE software, ACM transactions on software engineering methodologies, data centre world, empirical software engineering journal and data centre journal. The search was narrowed to only include publications for the period 1986 to 2016, which is 30 years of data aggregation, higher than the recommended 10-20 years. This was necessary as traditional air cooling (forced convection air cooling) has been used since the 1980's (McCartney, 1999).

Inclusion and exclusion criteria

The criteria used to include or exclude literatures are listed below.

1. One of the main inclusion criteria is that the papers must have undergone a thorough peer-review process, this is done to ensure the quality of the publications

used in this thesis.

2. Another criterion used for inclusion was to include papers which had an Informal approach to literature review, i.e. no defined research questions and no defined search processes, but they must enrich the ideas, thoughts or contributions made in this thesis.

3. Papers were excluded if they reported duplicates of the same study, including only the most complete paper.

4. Papers were excluded if they did not meet the Quality Assessment (QA) guidelines explained below in the quality assessment section.

Quality assessment

The QA criteria were formed by using the guidelines outlined by Keele (2007), papers were excluded from analysis if they did not meet any of the criteria listed below.

1. Do the papers provide any value for analysing cooling in LSDC's? (Value maybe anything which provides an insight to data centre cooling)

2. The results must have been validated either mathematically, or through experimentation, ensuring the published data can be relied upon.

3. Inclusion of white papers must not be biased towards a particular manufacturer, and must be based solely on the value these papers offer.

2.1.4 Results

An automated search on Google scholar using keywords like 'Liquid cooling in large scale data centres' returned 159,000 papers. 'Efficiency of liquid cooling data centres' supplied 45,500 papers and finally, 'Cooling requirements for next generation data centres' presented 47,000 papers. These keywords yielded the majority of the papers used, which is why the results of these are explicitly mentioned. Using other keywords provided additional papers, but are not mentioned here as the list would become protracted. The Google scholar searches were supplemented by referencing several white papers for emerging technologies like adsorption chillers and on-chip direct liquid cooling. For the full list of keywords used please refer Appendix A.2.

The industry architecture standard which applied to our research was the IEC 61346, which was referenced and adopted. The automated search resulted in approximately 252,000 relevant published papers among which around 100 qualified for primary analysis. Additional white papers and secondary references added another 25 papers. From this process we were able to gather 125 citations, from various sources to answer our research questions, all of which can be found in the 'References' section at the end of this thesis.

2.2 The challenges faced in LSDC's

2.2.1 The carbon footprint challenge

Data centres consume a lot of energy for cooling and system operation, which is why in 2006 all of the U.S data centres utilised 61 TerraWatt-hour (TWh) of electricity. This enormous amount of power consumption cost data centre operators an immense $4.5 billion (R. Brown et al., 2008). Sims, Rogner and Gregory (2003) estimate an average of 151g of greenhouse gases are released for every KWh of electric power consumed. This equates to 9,211,000,000 Kg or 9.211 Million metric tons of greenhouse gas emissions per year from the data centre industry in USA alone.

Electricity used in data centres mostly came from fossil fuelled sources, which is why the data centre industry is being heavily targeted to reduce their greenhouse gas emissions not only in USA but all over the world. David Cameron, the former UK prime minister introduced mandatory laws for data centre operators, among others to report their carbon emissions with the introduction of Carbon Reduction Commitment policies (Conservative & coalition government of the UK, 2016). These policies aim to tax companies which pollute the most and with talks of introducing a similar form of legislation in USA (Sultan, 2010), data centre operators are looking for ways to reduce their carbon footprints.

Global warming is a key driver in environmental policies and with all of these new legislations combined with rising power costs they pose a real world need for better solutions in reducing carbon footprints (Uddin & Rahman, 2011). Server life cycles exacerbate the problem of reducing carbon footprints, as they are changed frequently at the end of their life cycles or when it makes financial sense to replace them with more modern servers, which maybe as short as 2 years (Webb, 2008).

Data centre design challenge

Data centre designs take a minimum of 2 years to come into operation by which time they need to support the next-generation of servers. This is why green initiatives are being incorporated into data centre designs from inception (Murugesan, 2008).

Recent trends indicate that several next generation data centre operators are aiming to achieve carbon neutrality (Google, 2016b; Facebook, 2010). This drive to reduce greenhouse gas emissions is in part due to the increased awareness of green computing and the green revolution taking place across all industries (Webb, 2008).

Data centre design has been drastically changing with the introduction of modelling software to better accommodate new design changes for improving the air flow for cooling servers (Buyya, Ranjan & Calheiros, 2009). This is with the use of Artificial Intelligence (AI) to cut data centre PUE by 15% (Gao, 2014). The next generation data centres are being designed inside out and from within the servers themselves to further improve efficiency (Geet, 2014).

Some of the large data centre operators like Google are expecting to achieve fully carbon neutral LSDC's by 2017 (Google, 2016b). The cost to operate carbon neutral data centres are significantly higher than traditional ones. Case in point is Google's new data centre which uses recycled waste water for its cooling (Google, 2017b), adding to the operating costs as the water needs to filtered and recycled before being let off. Even if these next generation data centres achieve full carbon neutrality they constitute only a small fraction of the data centre industry.

While all of these measures aim to significantly reduce data centre greenhouse gas emissions and carbon footprints, the increasing server rack power density demands are offsetting these reductions. This is because server rack power density demands are outpacing data centre designs which is due in part to the exponential increase in advancements of computing technologies. Designing cooling systems capable of handling these increased loads remains a key challenge.

Data centre operators are turning to optimise cooling systems to help reduce their carbon footprints (David & Schmidt, 2014). This is because they are the largest consumers of electricity next only to the servers themselves, and implementing efficient cooling systems helps to reduce both power consumption and increase server component life cycles as investigated by Kreeley and Coulton (n.d.).

2.2.2 The cost challenge

The cost of operating LSDC's has exploded. According to Patel et al. (2002) a one hundred thousand square foot data centre will cost $44 million a year to power five thousand servers with power densities of 10 kW per server rack. Furthermore, they state it will cost an additional $18 million to power the cooling infrastructure at a total cost of $62 million in power consumption alone.

A study performed by Baliga, Ayre, Hinton and Tucker (2011) showed that cloud computing might not always be the most environmentally friendly way to offer storage, software and processing services in terms of power consumption. This issue is compounded when data centres idle or run at low utilizations. This phenomenon was studied by G. Chen et al. (2008) who investigated the power consumptions in data centres and found that, idle servers consume about 66% of peak power or full load.

According to Greenberg, Hamilton, Maltz and Patel (2008), power consumption is only one aspect of the total cost of a data centre, Table 2.1 illustrates the total amortized cost of a typical data centre.

Amortized Cost (Approx.)	Components	Sub-Components
45%	Servers	CPU, memory, storage etc.
25%	Infrastructure	Building, power distribution, cooling etc.
15%	Power Draw	Electrical utility costs
15%	Network	Links, transit, equipment etc.

Table 2.1: Amortized cost of a typical data centre

Interestingly enough the two major costs in a data centre, which are the server costs and the infrastructure costs are correlated. Romadhon, Ali, Mahdzir and Abakr (2009) asserts that proper cooling of IT equipment is directly dependent on its lifespan. However, this statement is argued by Webb (2008) who comments that the semiconductor industry has advanced to allow for greater tolerances.

Webb's statement is not entirely accurate. Their premise is that servers will retire earlier than they would fail by running them at hotter temperatures. This premise was evaluated by Kreeley and Coulton (n.d.) to determine a correlation between temperature and electrical equipment. They established that a hard drive would last 10 years when operated in a temperature range of 40 $^{\circ}C$ to 42 $^{\circ}C$. Furthermore, they found that when the hard drives operated at hotter temperatures their lifespans decreased significantly to nearly a year at 70 $^{\circ}C$.

This ambiguity in operating temperatures led the American Society of Heating, Refrigerating and Air Conditioning Engineers (ASHRAE) to stipulate operating guidelines for all the different classes of data centres (ASHRAE, 2011). These operating guidelines are followed by all data centre operators, where the maximum recommended operating temperature is 27 $^{\circ}C$. Data centre classes are based on the type of data, reliability, industry etc. to determine the class. As an example, an A1 class data centre would be a hosting enterprise grade infrastructure and a class A4 data centre would be hosting volume servers (ASHRAE, 2011).

Maintaining optimal operating temperatures is quite difficult when using air cooling systems. This is because air cooling systems are sensitive to temperature variations and are slow to respond to temperature fluctuations. According to Shrivastava et al. (2006). Air cooling IT equipment leads to 'hot spots', which are a huge problem in current air cooled LSDC's. Hot spots are areas of higher temperatures when compared to surrounding lower temperature zones. These hot spots increase the size and amount of cooling infrastructure needed to cool servers, thereby increasing the cooling costs for LSDC's.

2.2.3 Server rack power density challenge

Increased power densities offer greater performance but at the cost of unprecedented power consumption and heat loads at both rack and facility levels (Felter et al., 2003). Furthermore, Capozzoli and Primiceri (2015) explains the main reason for high thermal loads is the fact that IT equipment convert around 99-100% of the power they consume into heat.

Koomey (2011) asserts that the power consumption of data centres have roughly doubled between 2000-2005, and has continued to increase 56% between 2005-2010. This rapid increase in power consumption has placed heavy stresses on cooling systems which are struggling to meet the cooling demand placed by large heat loads. The problem of cooling is exacerbated when LSDC's employ hundreds or thousands of these high density server racks and place them in close proximity to each other. This is done to make operating LSDC's financially feasible (Jinkyun et al., 2009).

Traditionally LSDC operators used dedicated cooling equipment such as CRAC's to chill air and used server fans to regulate the temperature of servers (Jinkyun et al., 2012). This approach mandates the use of complex thermal management systems to regulate the operating temperatures of both the facility and server racks at a safe level (Chu, 2003).

Analysing the IT heat trend data for six technologies published by Romadhon et al. (2009), and inspecting future heat trends they have predicted. We can conclude optimising cooling in data centres has an excellent research scope.

Optimising cooling in data centres therefore is an excellent research opportunity and a key area for contribution. This is because current cooling systems account for around 40% of the total power consumption in a data centre (Capozzoli & Primiceri, 2015). Ebrahimi et al. (2014) have estimated that the power consumption of cooling systems in data centres are set to increase 15-20% annually. Power densities of server racks have increased from 1 kW per server rack (Patel et al., 2002) to future LSDC's being designed to operate at 100 kW per server rack (Musilli & Ellison, 2012).

2.2.4 The cooling challenge

As explained in the book 'Introduction to Heat Transfer' written by Bergman and Incropera (2011), heat transfer can occur through one of three modes, namely

1. **Conduction**: The process of heat transfer that occurs in a thermally conductive material or medium due to a temperature gradient. This states that conduction occurs through a material, for example in pots and pans whilst cooking. It is

interesting to note that this is the most efficient form of heat transfer.

2. **Convection**: The process of heat transfer that occurs between a fluid in motion and the surface of a material, if there exists a temperature gradient. The process is classified based on nature of fluid flow into either forced or natural convection. Natural convection is also referred to as free convection. We can determine the type of convection by calculating the dominance of diffusivity of molecules in fluid flows. If momentum diffusivity dominates then it can be classified as forced convection.

3. **Radiation**: The process of heat transfer that occurs due to the agitation of molecules. It mainly occurs through the emission of electromagnetic waves and the consequent interactions it induces on molecules.

The problem of cooling and the challenge of eliminating 'hot spots' using convection i.e. air cooling, is evident when a Computational Fluid Dynamics (CFD) analysis is done on LSDC's. Shrivastava et al. (2006) studied these hot spots and determined that, in some cases, these areas operated at temperatures far beyond the 'allowable' levels specified by ASHRAE (2011).

Next-generation air cooling systems tried reducing temperatures by adding raised floor ventilation systems and segregating hot and cold isles (Dai et al., 2014). Even with these optimisations, next generation air cooling systems have not managed to eliminate the phenomenon of hot spots. This is because air as a heat transfer medium is severely limited, it has a low specific heat capacity (the amount of heat required to raise the temperature of a given mass of substance) at 1.01 KJ/Kg when compared to a liquid such as water. Water has the highest specific heat capacity of any natural substance at 4.179 KJ/Kg. This makes liquid cooling ideal for high heat load applications, which is why it is used as the primary coolant even in modern nuclear reactors (Sinha & Kakodkar, 2006).

LSDC's have seen an accelerated increase in density, size and number (Barroso et al., 2013), this is because cloud computing is being adopted by all manners of industries like scientific, healthcare, education, consumer electronics etc. (Kearns, 2010). Trends indicate that the demand for cloud services will keep increasing with each passing year, which is why the scale of current data centres are increasing at an accelerated pace, both in terms of server rack densities and size, as investigated by Patel and Shah (2005).

The problem of air cooling is further compounded in emergency situations especially in high heat load applications. This is because when air cooling systems fail the time it takes before servers fail known as the 'Thermal ride-through' time is modest when compared to liquid cooled servers (Moss, 2011; 3M, 2016).

The thermal load balancing challenge

Air cooling is sensitive to thermal fluctuations as noted by Shrivastava et al. (2006). This leads to reduced efficiency and as such thermal load balancing should reduce these temperature variances. Sharma, Bash, Patel, Friedrich and Chase (2005) propose two methods of IT load distribution, as server temperatures are a function of IT load. They propose row-wise and row-row distribution techniques to optimise IT load distribution. These methods segregate the servers into isolated zones, this depends on the data centre topology.

Data centres use rows of server racks and numerous blade servers placed in these racks. The server racks have been standardised to allow for inter-compatibility between IT components. A study done by Patankar (2010) shows that the topology of these server racks has a huge impact on the air flow in air cooled systems. In liquid cooled systems, the topology between the server racks has a negligible impact but the topology of the connections between the server racks has a significant impact (Sickinger, Van Geet & Ravenscroft, 2014). Server racks can host a maximum of 42 1U rack mount servers commonly referred to as 'units'. A single unit (1U) is the size of a single blade server, a commonly used type of server in LSDC's (Pakbaznia & Pedram, 2009). The schematic of a 1U rack mount server is illustrated in later chapters of this thesis (refer Figure 5.1 in Chapter 5).

This method is limited in its application and is reserved for air cooled data centres. Mahapatra and Yuan (2010) explore network level load balancing mechanisms applicable for any data centre. These approaches are not implemented in commercial LSDC's thermal management system but rather proposed as a theoretical exercise. The limitation of this approaches is that it can reduce the phenomenon of hot spots but not eliminate it and as such hot spots prove to be a continued limitation in air cooling systems.

Uncontrolled thermal runaway problem

This thesis considers data centres as cyber physical systems with industrial automation controllers for thermal management of high density server racks. These thermal management systems use temperature sensors, actuators and other interface devices to avoid thermal runaway (Dwyer, Franklin & Campbell, 1990). Thermal runaway is a phenomenon where temperatures of components exceed safe operating guidelines, causing financial damage, and in some instances even loss of life. The thermal runaway phenomenon was studied by Wang et al. (2012) and found it was prevalent in lithium ion batteries, while Dwyer et al. (1990) found it to be prevalent in semi-conductor devices. This is why cooling systems are critical to maintaining safe operating conditions, which is especially true in high power applications such as data centres.

2.2.5 Impact of cooling on the challenges faced by LSDC's

Proper cooling impacts all of the challenges faced by LSDC's. This is because efficient cooling of data centre components increases their lifespans as shown by Kreeley and Coulton (n.d.). Cooling has a direct impact on operating costs as power consumption by cooling systems accounts for 40% of the total power consumed by data centres (Capozzoli & Primiceri, 2015). Lower power consumption has an indirect impact on the carbon footprints of LSDC's, which also affects data centre designs. Cooling systems have an immediate impact on the limit of power densities in server racks.

2.3 Methods of cooling data centres

2.3.1 Air cooling

The first generation of data centres had primitive air cooling technology, which were part of the air conditioning systems installed in large office buildings. Later when power costs began to rise, they used Computer Room Air Conditioners (CRAC) to make air cooling more efficient (Arguello-Serrano & Velez-Reyes, 1999). The main problem with air cooling is the formation of hot spots, which indicate high temperatures (Shrivastava et al., 2006).

Next generation air cooling systems reduces these hot spots by installing raised floors and perforated tiles to improve air flow within the server racks (Jinkyun et al., 2009). These steps reduced hot spots but did not eliminate them, which is why hot aisle containment insulation was done. This further reduced the number of hot spots, but still did not eliminate them (Madhavi, 2014).

The final enhancements made to improve cooling were to put in excess cooling capacity along with spacing the servers within the server rack far apart, i.e. using only 50% of the maximum server rack capacity (Shrivastava et al., 2006). All of these steps lead to increases in operating overhead costs, which is why the largest technology giants like Amazon, Facebook, Google, Intel, Microsoft, etc. have begun implementing alternate forms of cooling technologies (Askwig et al., n.d.; 3M, 2016; Facebook, 2010; Google, 2017b).

2.3.2 Hybrid cooling

Hybrid cooling technologies evolved from air-cooling, as air-cooling became too expensive to remain competitive. One of the hybrid cooling systems proposed was spray cooling (Lin & Ponnappan, 2003). The state of the art spray cooling systems takes advantage of the physical property of state change and the large amount of heat a state change requires. A state change is when water droplets in the humid air evaporate into water vapour. In tests conducted by Lin and Ponnappan (2003), they used water as the spray cooling method and were able to demonstrate spray cooling with heat flux of 500 watts per square cm for water. The sub-cooling characteristics for water (between 3 and 14.1 °C) were the greatest when compared with fluorocarbons.

These results are good enough for current LSDC applications which operate at power densities of 30 kW per server rack (R. Schmidt et al., 2005). But applying a spray hybrid cooling system to future LSDC's may lead to a humidity problem causing rain clouds to form inside the data centres. This leads to massive equipment failures if operating parameters are improperly maintained (Clark, 2013).

2.3.3 Liquid cooling

Liquid cooling as a concept is not new, but the implementation in data centre cooling is a new and emerging technology (Dai et al., 2014). Liquid cooling technology implemented in LSDC's looks promising with several technology giants researching the opportunity of liquid cooling like Intel (Nguyen et al., 2015) and Google already having implemented it (Gao, 2014). LSDC operators are levitating to liquid cooling systems because, water is 3467 times more efficient at removing heat than air (Aquatherm, 2015).

The enhancements made to improve liquid cooling efficiencies seek to leverage the advantages offered by legacy cooling technologies to implement liquid-air cooling technologies. This system uses a direct to chip liquid cooling heat sink for Central Processing Unit (CPU) and Graphics Processing Units (GPU) cooling and with other components cooled by air (Sickinger et al., 2014). This method takes advantage of the fact that most of the heat produced in a server is by the CPU and GPU chip and that other components in a server can operate at a much greater temperature.

2.4 Factors in choosing cooling systems

2.4.1 Data centre thermal topology

The topological configuration of data centre components plays a pivotal role in determining the method of cooling systems employed in LSDC's. This is because of the thermodynamic characteristics involved in heat transfer. As mentioned in the cooling challenge faced by data centres (refer Section 2.2.4), heat transfer is heavily dependent on the medium of heat conduction used, such as air or water. Heat transfer is also heavily dependent on the temperature differences between the mediums (Bergman & Incropera, 2011). Simply put the greater the temperature differences the faster the heat transfer.

Data centre cooling systems aim to remove the heat as fast as possible since servers are sensitive to large temperature fluctuations (Autodesk, 2013). Due to this sensitivity air cooling systems maintain lower air temperatures, which increases power consumption and leads to inefficient data centre designs being optimised for temperature differences rather than optimal power efficiencies.

In liquid cooled LSDC's thermal topological configurations illustrate the fluid dynamic relationships between the different data centre and server components (refer Figures 5.1 and 5.2). Liquid cooled LSDC's are especially sensitive to thermal topological configurations (hereafter referred to as data centre topology or simply topology), because the most commonly used connection in such systems is daisy chaining (refer Chapter 1.3.5). As the number of the servers in the daisy chain increases, the rise in temperature of the coolant also increases, thereby decreasing the rate of heat transfer. This requires increased coolant flow which decreases power efficiencies.

Factors in choosing liquid cooling technologies

Liquid cooling technologies can be differentiated based on the state change (refer Section 2.4.2) properties of the coolant fluid. State change refers to the conversion of a liquid coolant into a gaseous state, which occurs during heat absorption. The different types of liquid cooling are listed below,

1. Single phase: Where the coolant does not undergo any state changes
2. 2-Phase: Where the coolant undergoes a liquid-vapour state change by evaporation.

Single phase: Single phase liquid cooled systems have been used in various industries for high heat load applications, for example, mainframes shipped with single phase liquid cooling systems (Simons, 1995). The advantage of single phase liquid cooling systems is that they are simple to implement and require relatively less complex thermal management systems.

2-Phase: An emerging technology in liquid cooling systems is 2-phase immersion liquid cooling, which has been tested thoroughly by bitFury group and found out to perform other cooling methods with a PUE of 1.01 (3M, 2016), which is nearly the theoretical limit (1.0). The coolant used for the ground breaking 40+ MW data centre being constructed in the republic of Georgia is 3M's Novec 7100 fluid, which is a 2 phase immersion cooling liquid with low global warming potential.

A 2-phase system would be ideal for our research as it is cutting edge, however, 2-phase liquid cooling systems are incredibly complex and would lead to greater financial losses in the case of emergencies, like the loss of coolant flow. Other reasons for opting a single phase solution instead of using a 2-phase system is because a single phase system requires less infrastructure such as the lack of server containment units and hermetic sealing of server units or the presence of expensive condensers inside server units, which are necessary in 2-phase systems (Tuma, 2010).

Another key design choice for our system is that we intend to develop systems which possess the capability of being retro fitted to existing current generation LSDC's. The 2-phase liquid cooling system needs complex infrastructure which can be only implemented in future LSDC's.

2.4.2 Limitations of cooling systems

The limitations of air cooling systems have been thoroughly discussed in the introduction (see Section 1.3 and 2.2). The prevalence of air cooling systems can be explained in part due to the leakage problem faced by implementing liquid cooling systems. Leakage of coolant occurs when transporting the coolant to and from various parts of the cooling system. The problem of leakage is accelerated if the coolant heats up beyond the operating temperature of the fluid, when the fluid turns to gas and adds pressure to the coolant pipes causing leaks. With new innovations in fluid dynamics, the leakage problem has been eliminated (Aquatherm, 2015).

Leakage problems were not evidenced when Zeighami, Saunders, Coles and Branton (2014) conducted thermal tests to determine if waste grade heat energy could be captured and reused. Solutions offered by Asetek were a really good fit for our research as our contributions can be applied to their models. There are a few drawbacks to this solution, as CRAC's are still needed, along with the infrastructure for liquid cooling which may decrease the financial incentives for adoption in existing current generation LSDC's.

These limitations are dwarfed when compared to the limitations of hybrid cooling systems which require complex CRAC and HVAC management systems. Hybrid systems are also prone to risks of corrosion and pre-mature equipment failure. Failures can occur when humidity levels are maintained incorrectly, which cause corrosion in various data centre components, which makes them infeasible if implemented or maintained improperly (Klein et al., 2011).

Hybrid cooling can cool current generation LSDC, but next generation LSDC cooling is yet to be conclusively proven. The corrosion effect occurs when water is used but using other coolants such as fluorocarbons affect the environment heavily. Therefore, hybrid cooling systems applicability in future LSDC's is limited.

Taking a pragmatic approach and critically analysing the limitations of all the various cooling systems we found that liquid cooling offered the least number of technical challenges, which could hinder adoption in current and next generation LSDC's.

2.5 Cooling control strategies for liquid cooled LSDC's

Control strategies for liquid cooling systems can be differentiated based on the nature of coolant flow, namely

1. Static flow control strategy and
2. Variable flow control strategy.

2.5.1 Static flow control strategy

According to Steinke and Kandlikar (2006) the flow of coolant remains constant throughout their operation. A static flow control strategy is employed when the thermal dissipation of the system is known and is likely to be minimal as was the case in legacy mainframe computers. Mainframes had high thermal dissipations with minimal fluctuations, which is why they shipped with statically controlled liquid cooling technologies

(Simons, 1995). Static flow liquid cooling has also been used in advanced systems, like the IBM 775 supercomputer (Ellsworth et al., 2012), as the system's heat output was constant and did not vary enough to allow for fluctuations in coolant flow. This made implementation of static flow cooling systems feasible.

Modern LSDC's have a wide gamut of temperature ranges, which fluctuate based on the IT load placed on them. This uncertainty in thermal dissipation reduces the efficiency of statically controlled liquid cooled LSDC's.

Limitation of dynamic cooling in statically controlled systems

The main challenge in using a static cooling control strategy in LSDC's is that it leads to over-cooling or under-cooling of servers. This is because as explained earlier, server temperatures are heavily dependent on the IT load placed on them and since the IT load varies significantly with usage demands the heat profiles of the servers vary constantly. Using a static control strategy offers very little flexibility to optimise the coolant flow to meet these variations. Over-cooling is a problem as it consumes more power and under-cooling leads to pre-mature equipment failure and financial losses through server downtime. To mitigate these problems dynamic cooling was employed to allow for coolant flow to synchronise with server cooling requirements.

2.5.2 Variable flow rate cooling control strategy

In these systems, the flow of coolant can be varied at any time. This is done as a method of dynamically responding to changes in the IT load or to respond to dynamic variation in temperatures (Marcinichen, Wu, Paredes, Thome & Michel, 2014). Variable flow liquid cooling systems fluctuate the coolant flow continuously based on various criteria, which is why they require advanced management systems such as the AI system used by Google (Gao, 2014).

Variable flow cooling systems have been evaluated by Marcinichen et al. (2014) and found to be a more robust and flexible system, which is why they have been implemented in Google's LSDC's (Gao, 2014).

2.6 Findings

2.6.1 Gaps and opportunities

We have comprehensively assessed the impact of cooling on the challenges faced by LSDC's (refer Section 2.2.5) and found that it has a direct cause and effect relation for server equipment lifespans, power consumption, data centre designs and most importantly carbon footprints. We have shown data centre thermal management systems are incredibly complex and use AI and/or automation controllers to maintain safe operating temperatures (Google, 2017a). We have also voiced that data centres are cyber physical systems which makes it possible to implement complicated AI cooling control systems for thermal management. The limitation of using AI according to Gao (2014), is that the data collected from sensors needs to be accurate and processed in real time for the system to function effectively.

The AI system implemented by Google (Google, 2017a) uses a continuous variable flow rate, where the flow rate of the coolant is continuously adjusted to meet the servers cooling demands (VFR cooling control strategy). Another limitation identified with Google's implementation is that it requires large amounts of computing resources since the sensor data needs to be cleaned and processed in real time, with control logic being applied to this data and then varying the flow rate to meet the cooling demand. Such a long feedback loop adds to complexity and chances for errors to creep into processes.

These limitations in current liquid cooling systems and control strategies provide us with a solid foundation for further research and contribution to be made in the field. The focus of this thesis is to optimise liquid cooling control strategies such that they conform to the cooling requirements of current and future generation LSDC's. Liquid cooling technologies were chosen as they provide the greatest opportunities for contributions. This is because liquid cooling technologies have not matured (Dai et al., 2014) which provide valuable gaps in knowledge. These gaps offer research opportunities which can be leveraged to propose solutions and make contributions in the field of cooling in LSDC's.

2.6.2 Addressing research question 1

The data obtained from the SLR process yielded the answer to our first research question i.e. What are the key functional and non-functional requirements for cooling systems in next generation LSDC's?

This research question is thoroughly answered in later chapters of this thesis (refer chapter 4), but a summary of a few important aspects are listed below,

Cooling requirements

Next generation data centres utilise power densities of at least 100 kW per server rack (Musilli & Ellison, 2012). Data centres are required to be environmentally friendly, making these carbon neutral in operation can be achieved through reducing power consumption and implementing innovative technologies (Google, 2016b, 2017b). Cooling systems are targeted for power reduction as they are the most power consuming equipment after powering the servers themselves. Cooling systems consume around 40% of total power (Capozzoli & Primiceri, 2015) in data centres while the rest is used to power servers and data centre facilities like lighting. This is why research has been ongoing to determine suitable forms of alternate cooling technologies and their feasibility (3M, 2016; Madhavi, 2014).

Topology and data centre designs

Topology and its impact also affects data centre designs, which is why the physical footprints of data centres are decreasing, correlating to the increase in power densities of server racks (Uddin & Rahman, 2011). Designing green data centres is a key requirement for future LSDC's (Geet, 2014), hence the power efficient strategies involved should be able to comply with different data centre topological configurations.

2.6.3 Optimising liquid cooling control strategies

We hypothesise that feeding pulses of power rather than using continuous variable flow of the coolant leads to better efficiency. This can be achieved through Pulse Width Modulation (PWM) (Gobor, 2016). We also aim to use shorter feedback loops by deploying the control logic on programmable logic controllers (PLC's) to reduce the amount of computing resources required, a limitation posed by Google's AI system.

Our hypothesis is based on three main principles, namely

1. The efficiency of coolant pumps is maximum at near full load (75-80%) (Cho et al., 2007).

2. Feeding pulses of power instead of supplying continuous power decreases power consumption (Choi, Kong & Choi, 1994).

3. Deploying control logic on PLC's reduces the computing resources needed for execution (Dubinin & Vyatkin, 2007).

The main difference between the continuous Variable Flow Rate (VFR) cooling control strategy and Pulsed Variable Flow Rate (PVFR) cooling control strategies is that we intend to maximise the cumulative efficiencies of all component devices in the network. This means to say that, by optimising and increasing the operational efficiencies of all the small insignificant aspects involved in cooling, we will observe a net decrease in power consumption. The PVFR control strategy aims to keep the pumps off as much as possible, whereas the VFR strategy is to keep pumps on as much as possible. This contrast is against conventional wisdom, and the effectiveness or efficacy is analysed and discussed later in this thesis (refer Chapter 6).

Optimisations are based on the iterative framework proposed by model driven design (refer Chapter 5). This is done to ensure a significant power reduction is obtained which allows for greater adoption by existing and future LSDC operators.

Assumptions

We have made a few assumptions in our hypothesis, we have assumed the net cooling effect of a pump from idle to full power is the same in all aspects. This is because modern starters for coolant pumps have eliminated the spikes in power consumptions that used to occur in legacy pumps (Goh, Looi & Kok, 2009). While the power consumptions have improved, the latency between the two control strategies is still present. This difference is less pronounced in liquid cooling systems, as reservoirs and micro-channels (Rezania, Rosendahl & Andreasen, 2012) eliminate any ripples in coolant flow and smooth the flow of coolant over heat sinks.

Another assumption made is that the sine-wave (mains power) to PWM wave power conversion losses are low, this is because high efficiency converters are prevalently used (Wai, Lin, Duan & Chang, 2008). Wai et al. assert that the conversion efficiency is in excess of 91%. This conversion is needed only when power is being consumed, and in the case of PVFR the losses are negligible when compared to the continuously varying, VFR control strategy.

2.6.4 Testing the hypothesis

Our hypothesis is that using pulses of intermediate coolant flow rather than employing continuously varying or static coolant flow will reduce power consumption of the cooling systems whilst maintaining safe operating parameters (refer Section 1.5). To properly test this hypothesis, we need to implement the different cooling control strategies on physical data centres. This approach is not practically and financially feasible as the amount of money required to gather the IT components and power them will far exceed any available sources of funding.

Another reason to support using computer simulations rather than physical data centres is that data centre topologies vary considerably for different regions (Jinkyun et al., 2012). This is because data centres are optimised for given environmental conditions. Variability of data centre topologies makes extrapolating the data gathered from them difficult to be applicable for other data centre topologies.

These challenges make implementing different control strategies on physical LSDC's currently infeasible. Which is why computer simulation software's were chosen to test the various different control strategies needed to develop an optimised cooling control strategy. The focus of optimisation is to reduce the power consumption of coolant pumps in LSDC's. This is because cooling systems account for 40% of the total power consumed in LSDC's (Capozzoli & Primiceri, 2015) and offer the greatest opportunity for contribution (refer Section 2.6).

Current simulation software like CloudSim are only capable of simulating air cooling technologies. Another major limitation is that they can only simulate the power consumptions of CRAC's, HVAC's and chillers, which are the components used in air cooling technologies and not liquid cooling systems (Buyya et al., 2009). While SimScale (Scale, 2017) offers a more comprehensive set of thermodynamic CFD simulations, but fails to simulate power consumptions of coolant pumps and the interactions of data centre topologies on the scale of LSDC's.

Autodesk (2013), one of the leaders in CFD simulation software also fail to meet such simulation requirements. This led us to develop a novel closed loop single phase simulation software known as DynaCool. DynaCool is capable of simulating the Dynamic cooling control strategies hence the name. In its current form, it is limited in terms of CFD simulation capabilities, in that it focuses only on closed loop single phase simulations. A more comprehensive list of the limitations of the DynaCool system is identified and left as opportunities for future researchers which are explained in later chapters of this thesis (refer Section 7.3.4).

Due to these limiting factors we intend to develop and simulate using the proprietary DynaCool simulation software (refer Chapter 5). The simulation models are based on the real world models proposed by Marcinichen et al. (2014), Autodesk (2013), Musilli and Ellison (2012) and Kang, Miller and Cennamo (2007). Before gathering data and inferring any results the DynaCool system needs to undergo extensive validation against real world physical data of LSDC's (refer Section 5.2.3).

The need for a Requirements Engineering (RE) process to develop DynaCool

Bell and Thayer (1976) asserts that if the prevailing school of thought that requirements just arise naturally is followed, it might lead to total software failure even though it might be technically correct to do so. They go on to theorize that if coding personnel are given a needs statement for software implementation, which is a critical phase of software design they might be overlooked. The lack of a thorough system design might make the software fail to meet even basic functional criteria. Which is why for designing the DynaCool system we use modelling languages and a systematic approach for requirements elicitation.

Alford (1977) dismisses this theory entirely and adopts the notion that requirements are merely something that's needed and failing to implement some of the requirements will not lead to such catastrophic results. Adams (2015) simplifies the classification of requirements by classifying them into two paradigms, functional and non-functional. The word paradigm ensures that they are concerned with separate aspects of software development.

Functional requirements according to Kossiakoff, Sweet, Seymour and Biemer (2011) are largely concerned with what the system should do, and these requirements govern the systems actions while usually being action oriented. They accurately describe the operational tasks which the system performs. The essential characteristics of functional requirements according to Adams (2015) are that it should explicitly define what the system does, be action oriented, describe tasks or activities and be affiliated with input to output transformation.

It would be beneficial to understand what are the characteristics of good software requirements. Committee and Board (1998) provides guidelines for formulating good requirements which are, a good software requirement should be *"Correct, Unambiguous, Complete, Consistent, Ranked for importance or stability, Verifiable, Modifiable and Traceable"*. Using these characteristics, we can formulate proper requirements for our DynaCool system (refer Section 4.1.3).

2.7 Conclusion

Data gathered from the SLR process (refer Section 2.1) allowed us to address the first research question. This analysis identified an urgent need for better, more efficient cooling systems for current and next generation LSDC's. Market validation was explored by analysing the financial and environmental savings offered by liquid cooled systems. Research and implementation for alternate cooling technologies by various LSDC operators such as Google, Facebook, Intel etc. (Google, 2016a; Facebook, 2010; 3M, 2016; Askwig et al., n.d.) proved the viability of exploring the problem of optimising cooling in LSDC's. This problem of optimising liquid cooling control strategies is our second research question, which was answered using model driven design and engineering (refer Chapter 4).

Data suggests that reducing PUE by even a fraction of a percent saves millions of dollars in operating costs for LSDC's. Our focus in this thesis is to optimise cooling efficiencies by employing a novel PVFR approach. This approach aims to reduce PUE by reducing the power consumption of coolant pumps which are critical in liquid cooling systems.

Chapter 3

Research Methodology

This chapter details the methodologies used to undertake the research this thesis, Section 3.1 illustrates the selection of a suitable research methodology for our application. Section 3.2 discusses the analysis of different development methodologies used for our system design. Section 3.3 asserts the validation methodologies used to ensure the correctness of the system being designed and developed. Finally, Section 3.4 provides a summary to the chapter.

3.1 Selection of a suitable research methodology

Researchers often have a strong tendency to associate their work into two paradigms, qualitative or quantitative methodologies; and in doing so accentuate the apparent different philosophical roots. Kelle and Erzberger (2004) assert that there need not be an impenetrable frontier among the two. They argue that recent trends have blurred the divide, as an example, focus groups and non-standardized interviews are being linked to quantitative methods and is leading to joint research design practices. Kaplan and Duchon (1988) and Morgan (2007) support the same premise of integrating both methodologies. Kaplan and Duchon (1988), explain that this should be the norm, especially the case in information systems. While they do make a good case, determining an accurate methodology is hard, as supported by Easterbrook, Singer, Storey and Damian (2008) who suggest that well documented literature of the pros and cons for any particular method is not always available.

3.1.1 Candidate methodologies

Easterbrook et al. (2008) identifies several approaches that were considered and examined for viability while evaluating the DynaCool simulation software. The methodologies that were not viable are listed below as alternatives that were considered before being eliminated.

Action research

This methodology combines theory and real world use case scenarios, it is a quantitative methodology which is mainly used to explain the behaviour of software from an organizational perspective (Avison, Lau, Myers & Nielsen, 1999).

Ethnography research

This methodology assembles researchers from perspective fields in a particular subject area to examine all available data, and is similar to action research in this regard. Both action and ethnography research methods draw conclusions arising from observations and outlying data that might not have been reported or seen as relevant unlike other methodologies.

Qualitative methods involving case studies and surveys

Case studies and surveys have proven to be effective given that they provide enough data to draw tangible results (Kelle & Erzberger, 2004). Developing models for DynaCool could prove helpful if data is obtained from industrial experts. This was considered and several data centre architects were approached and surveyed, but there was insufficient data to make any meaningful design decisions.

3.1.2 Our approach

Considering all of the research methodologies listed above, in this instance, our research gravitates more towards a quantitative approach that includes data gathering and analyses. Which is why we have adopted quantitative methods along with experimentation and statistical analysis for our research. Most researchers in software engineering, document experiments badly. Which was evident when Sjøberg et al. (2005) surveyed 5453 articles from leading journals. Who along with Dybå, Kampenes and Sjøberg (2006) concluded that most researchers and the experimental methodologies they used were documented inadequately and had used extremely poor experimental design. Hence, several authors such as B. Kitchenham et al. (2008) Jedlitschka, Ciolkowski and Pfahl (2008), and Jedlitschka and Pfahl (2005), have proposed the following minimum discussions that should be published in any study.

1. Related work, to understand the domain of knowledge (refer Chapter 2)
2. Experimental design, describes the outcomes and hypothesis used in the experiments among other things (refer Chapter 6).
3. Execution, to understand how the design was implemented (refer Section 6.2).
4. Analysis, summarising the data and to describe how the data was analysed (refer Section 6.3.3).
5. Interpretation, to draw correlations from the data if any (refer Section 6.4).

6. Conclusions and future work, summary of the whole study and potential opportunities for future researchers (refer Chapter 7).

3.1.3 Systematic Literature Review (SLR)

SLR as a method for analysing literature in the domain has been explained previously in the literature review chapter of this thesis (refer Chapter 2). Purpose of this discussion is to critically evaluate the research problem and identify the limitations of existing solutions. Our research domain as alluded to in the introduction (refer Section 1.2) is one of cooling in LSDC's and is a vast area to cover which is why using a SLR (refer Chapter 2.1) helps to avoid the inevitability of limited scope which ensue in other review methods.

3.1.4 Design methodology

Purpose

Before designing experiments, we must first understand the role of statistics in science and engineering. Consider determining the probability of heads or tails in a coin toss. If we flipped a coin ten times and the outcome was there were six heads and four tails, does this mean that we can infer the probability of heads is greater than that of tails? If we used a purely empirical method to evaluate the data, then this is what would be inferred as factual. But the fact is that there is equal probability for either heads or tails.

There is a certain level of uncertainty in the universe which is especially true in thermodynamics (Zhu, Ren & Li, 2009). Thus a purely empirical approach will not suffice and this is why statistical experimental design is used to eliminate such ambiguity and biases that may be introduced knowingly or unknowingly (Box, Hunter, Hunter et al., 1978) and (Chandler, 1987). It has been well established that statistically designed experiments are used to eliminate sources of personal biases. Which is why they are used to ensure precision, providing reliable inferences which can be reproduced through experimentation (Mason, Gunst & Hess, 2003).

Framework

Mason et al. proposes a framework for practitioners, which consist of three distinct phases, these relate to the measurement of variation in sources. This is similar to the four phase framework proposed by Basili, Selby and Hutchens (1986). The only difference between the two is a definition phase, which provides an introduction to the study. The definition phase is incorporated into the planning phase by the former proposal. All frameworks have essentially the same critical stages of problem definition, experimental planning, operation or execution and interpretation or statistical analysis according to Mason et al. (2003).

The motivation and purpose of study has been explained in the introduction and literature review chapters respectively (refer Sections 1.1 and 2.1.1) along with the domain and focus of study (refer Sections 1.2 and 2.1.2). Which led us to define the problem in previous chapters (refer Section 2.1.2). To evaluate the effectiveness of the proposed PVFR control strategy we need to design experiments for gathering data.

The design of this study couple analytical methods with the scope and indicates the domain samples examined. Domain samples include existing cooling systems and their control strategies identified in previous chapters (refer Section 1.3.2). There are many statistical data analysis designs such as ANOVA, t-tests etc. (Box et al., 1978). However, for the sake of simplicity, accuracy and ease of analysis we will employ t-tests. This is because t-tests are accurate and precise for the data we gather from contiguous parametric measurements i.e. interval data to determine the efficacy or effectiveness of PVFR.

3.2 System design methodology

Hehenberger et al. (2016) assert that the main drivers for the design, modelling and development of software for Cyber Physical Systems (CPS's) are the reduction in development costs and time. One such method according to Gomaa (1989) is to use an Object Oriented (OO) approach to system design. While this is used heavily throughout the software industry, in real-time embedded control systems, an OO approach fails to account for interactions between tasks like the latencies encountered in control systems. While they propose a new system design approach its implementation has not been well established in the software industry (Gomaa, 1989).

Bachmann and Bass (2001) proposes using an architecture to account for uncertainties in system designs. This notion is also held by Hehenberger et al. (2016), they specify separating the various system layers into abstracted models and controller using formal modelling languages such as the widely adopted Unified Modelling Language (UML) among others. The UML diagrams and the system design for the software we use, called 'DynaCool' are detailed in later chapters of this thesis (refer Section 4.2.2).

3.2.1 Architecture selection

As we intend to abstract the cooling control strategies with the data centre plant model, we opt for a Model-View-Controller architectural pattern. This architecture offers abstraction and modelling of both data centre plant or facility and cooling control strategies. The reasoning behind this decision is explained thoroughly in later chapters of this thesis (refer Section 4.2.1). But as IEC 61499 is an industry standard predominantly used in embedded control systems, and has a few useful elements such as event driven execution, performance etc. we have incorporated such features into the final system design (refer Section 4.2).

3.2.2 Development of a solution

The system design requirements as alluded to earlier involve proper abstraction of models and controller, the complete set of requirements identified during the formal requirements engineering process has been detailed in later chapters (refer Chapter 4). To meet these requirements a design approach was considered that it is a robust method used in information systems (Peffers, Tuunanen, Rothenberger & Chatterjee, 2007). The main limitation of this approach however is that it is focussed heavily towards empirical evidence. Our research as explained earlier (refer Section 3.1.2) leans more towards statistical analysis rather than a purely qualitative based approach. This is why using a design science methodology would prove to be inapplicable.

We have opted to use a model driven design or Model Driven Engineering (MDE) approach for an iterative development of the DynaCool system. This is because the feasibility and accuracy of simulation software are very difficult to calibrate properly initially (refer Section 5.2.3). An MDE approach allows for proper calibration and development of both data centre plant/facility models and Controller which incorporate different cooling control strategies. The iterative development process is outlined in later chapters of this thesis (refer Section 5.1.3).

An MDE approach is most commonly used in mission critical iterative development of complex systems (France & Rumpe, 2007). While using an MDE approach is often time consuming and complex, cooling systems in LSDC's are mission critical and failure of such systems would prove to be catastrophic and as such the applicability of an MDE approach is justified.

3.3 Validation methodology

3.3.1 System validation methodology

Criteria for simulation software selection tend to be either direct reflections of cost and reliability (Wolverton, 1974; Littlewood & Verrall, 1973) or indirect reflections of complexity (McCabe, 1976) to name just a few of the elements in software science (Halstead, 1977). These criteria are devalued as we have developed a proprietary software known as DynaCool to simulate the experiments. This is done as we lack access to physical LSDC's, and mainly because we are predicting a 100 kW per server rack density for next generation LSDC's (Musilli & Ellison, 2012). These server racks presently do not exist and the availability of software's that can simulate liquid cooling of LSDC's are absent. Some similar software like CloudSim and GreenCloud do simulate power and offer CFD calculations, but are limited to air cooling and not liquid cooling.

Our goal is to analyse if a statistical significant power reduction can be achieved by using different implementations of dynamically liquid cooling systems. The factor criteria used is a p value less than 5% or $\alpha = 0.05$, this means that we can be 95% confident that the results are valid for all cases. This level of confidence offers us to infer a strong correlation and is generally the cut-off used in all statistical tests, the level of measurement is contiguous interval. We compare the data gathered using standard system validation tests through the DynaCool system with the data published by other researchers in physical data centres discussed in later chapters (refer Section 5.2.3).

3.3.2 Data validation methodology

Statistical data analysis methods are many and used across all scopes of experimentations, some of which are correlation, regression, factor analysis etc. (Neter, Wasserman

& Kutner, 1974; A. T. Allen, 1982). In this instance we are establishing a correlation between the power consumption of coolant pumps for different cooling control strategies. Which implies determining if implementing one cooling control strategy over others results in a statistically significant reduction in power consumption of coolant pumps. Sampling techniques maybe used to get a better selective sample, (Cochran, 2007) but this is not representative of real world use cases, and as such no sampling techniques are employed.

As explained previously (refer Section 3.1.4), many methods of statistical data analysis exist such as t-tests, chi-square, G and ANOVA tests (Mason et al., 2003). Many used in evolutionary studies to identify regression, apply a curve fit and to calculate probability measures (Box et al., 1978) and are utilised to predict future data based on past trends (Rice, 1989). The selection of a suitable data validation method is discussed in the analysis chapter of this thesis (refer Section 6.3.1).

3.3.3 Experimental validation methodology

We need to determine an optimal static flow rate which acts as comparison measure against which other dynamic flow rate implementations can be compared, which is our pilot study. A set of experiments need to be performed to get a spectrum of data for various conditions. This is done as it helps to mitigate biases introduced knowingly or unknowingly during experimentations. The minimum set of experiments needed are the efficacy, extended time, increased server rack density and increased number of server racks. These experiments should help us determine a true correlation for all use cases. These are the minimum set of experiments as they help to observe differences in prolonged usage, future use case scenarios i.e. scalability or in other higher thermal output applications and different data centre topological configurations. Data metrics and analyses can be performed by gathering data after completion from all the above experiments (Basili, Selby & Phillips, 1983).

3.3.4 Interpretation

The context is derived from the purpose of study, domain and the statistical framework to derive results (Basili & Reiter, 1981). Context is important as this avoids confusion while peer reviewing and makes the study valid in the research domain niche. Interpretation of the results and the validity of the data in the context is an important topic to clearly and explicitly state as the results are usually extrapolated by researchers to determine its applicability in the research area. An analysis of the sample usually gives a representative idea for extrapolation (Basili & Selby Jr, 1986). Impact of a study is usually determined by the conference or journal they are published in, along with the feedback offered by other researchers in the domain, its applicability in real world commercial systems and their replication potential (Box et al., 1978).

Figure 3.1 below illustrates the summary of the experimentation framework adopted for our study and this thesis.

3.4 Summary

Section 3.1 discussed the selection of a suitable research methodology and its application on our research, the nature of which was identified previously (refer Chapter 2). Section 3.2 outlined the methodologies used for designing the DynaCool system, which will be used later on while developing the system (refer Chapter 4). Section 3.3 detailed the exhaustive process of different validation methods used to ensure the data gathered and the results inferred are correct. These methods are used in subsequent Chapters (refer 5 and 6).

1. Definition					
Motivation	***Object***	***Purpose***	***Perspective***	***Domain***	***Scope***
To improve dynamic liquid cooling efficiency	Max server rack temp. and power consumed by cooling pumps	To evaluate the efficacy of varying implementations of dynamic liquid cooling over static liquid cooling	Researcher	Analysing total power consumed by cooling pumps	Single project

2. Planning		
Design	***Criteria***	***Measurement***
Experimental Design: Completely Randomised Control experimentation	Proprietary PowerSim software v2.1 for data collection	**Metric Definitions:** **Goal:** Statistical significance of power reduction **Factor Criteria:** $p < 0.05$
Multivariate Analysis: Correlation		**Level of Measurement:** Interval data
Parametric sampling is used		

3. Operation		
Preparation	***Execution***	***Analysis***
Pilot study will be performed to determine optimal static flow rate	Data is logged automatically by the software	Quantitative data analysis using preliminary techniques mapped using graphs

4. Interpretation		
Context	***Extrapolation***	***Impact***
Statistical framework	Analysis of sample for representativeness	Based on Real world application, published material and ability, easiness of the study to be replicated by peers during a review process.

Figure 3.1: Summary table of the experimentation framework

Chapter 4

DynaCool System Design

This chapter provides details for designing the DynaCool system such that it meets the cooling requirements for LSDC's. The organisation of this chapter is based on the formal Requirements Engineering (RE) process outlined by Kotonya and Sommerville (1998).

Section 4.1 discusses existing computer simulation software and their limitations, the stakeholders identified for a proprietary DynaCool system and the requirements elicited for designing the system. Section 4.2 outlines the model driven design process used for creation of valid models and the architecture employed to develop the system. The models created include both the data centre plant/facility and controller which incorporates the different cooling control strategies identified in previous chapters of this thesis (refer Section 2.5). Section 4.2 illustrates the UML (Use Case and Activity) diagrams needed for implementation of the system along with iterative development of the DynaCool controller. Finally, Section 4.3 summarises the system design of the DynaCool simulation software system.

4.1 System requirements specification

The Requirements Engineering (RE) process according to Sharp, Finkelstein and Galal (1999) involves stakeholder analysis, requirements elicitation, which incorporates both functional and non-functional aspects and finally requirements analysis. Requirements elicitation incorporates both functional and non-functional aspects and requirements analysis includes system development design. The RE process is applied for the design of DynaCool system as the development of this proprietary simulation software meets our simulation requirements as outlined in the literature review chapter (refer Section 2.6.4).

4.1.1 Stakeholder analysis for the DynaCool system

Data centre infrastructure designers use CFD simulation software to design data centres (Shrivastava et al., 2006). Data gathered from these computer simulations help data centre infrastructure designers optimise the design of data centres to enhance their efficiencies. This is the accepted practice while designing current and future LSDC's (Patankar, 2010; Zeighami et al., 2014). This makes data centre infrastructure designers a key stakeholder.

Data centre operators have a vested interest in minimising their operating overhead costs, and according to Parolini, Sinopoli, Krogh and Wang (2012) designing intelligent system controllers leads to significant energy savings in LSDC's. While the term 'intelligent' is ambiguous, it refers to embedding a control logic into the controllers. In this instance we intend to embed cooling control logic, this is because we are focusing only on optimising cooling as explained previously (refer Section 4.1). This evidence suggests that data centre operators and associated staff like analysts, strategists and administrators are stakeholders.

Researchers like us who are investigating innovative cooling control strategies for applications in all high thermal output systems can benefit from CFD simulation software. for example, in mining, medical, scientific and high technology industries. In this instance, researchers benefit from simulation software which are capable of testing different cooling control strategies for different data centre topologies. This makes researchers looking to optimise liquid cooling technologies stakeholders.

To summarise, the stakeholders we have identified and as such our focus for designing and developing the DynaCool system are listed below.

1. Data centre operators (Analysts, Strategists, Administrators etc.)
2. Data centre infrastructure designers and
3. Researchers looking to optimise liquid cooling technologies in data centres

4.1.2 Requirements elicitation

To gather a comprehensive list of requirements we need to elicit all the requirements targeting the above stakeholders in the applicable domain. To understand the domain, we use the following processes, which are based on the design, modelling and development guidelines proposed by Hehenberger et al. (2016).

1. Process modelling for understanding the topological fluid dynamic relations which interconnect the different components in a liquid cooled LSDC.
2. Prototyping various versions of DynaCool using model driven engineering to optimise and meet the quality attributes.
3. Document analysis for understanding the requirements for the system.

Process modelling

The topology of LSDC's have an immediate impact on the efficiencies of cooling control strategies as discussed in the literature review chapter (refer Section 2.4.1). The topology of LSDC's plays a crucial role in the design of the DynaCool system. The most common way that LSDC's connect servers and server racks is by using daisy chains. Daisy chaining is popular because of the modular nature of server racks (Hastings et al., 2002) and the ease of providing redundancy needed for server components. Redundancy is a critical part of operations and maintenance of data centres (Sanders et al., 2006).

In daisy chained liquid cooled server racks the coolant input to a server rack is the output of the previous server rack (Madhavi, 2014). In closed loop systems, this daisy chaining poses a problem as the coolant temperature will increase drastically after a few passes (Kang et al., 2007). This is why heat exchangers are used to cool the working fluid to maintain safe operating temperatures. Using traditional liquid to air heat exchangers proves to be ineffective at cooling the fluid in the required amount of time. The selection of the heat exchangers is based on the requirements outlined by Kakac, Liu and Pramuanjaroenkij (2012) and is modelled later in Section 5.1.2.

Model driven engineering

France and Rumpe (2007) asserts that model driven engineering is used to reduce the complexity during implementation of complex systems. Model driven engineering uses iterative development practices to refine and validate the models. This stepwise development allows to test and validate developed models early in the implementation time frame. In this instance, our DynaCool system design needs to be abstracted into two levels. This is because the computational network affects the thermal output of the server racks and implementing different cooling control strategies has various impacts on the thermodynamic characteristics of the thermal network.

Using an UML and MDE approach allows controllers to leverage interactions and characteristics of the various devices and networks in LSDC's. This allows for the implementation of distributed strategies ensuring greater reliability, which allows modularity among controllers to enforce programmed intelligence logic to maximize efficiency of cooling through collaboration and cooling control independence. However, this requires a distributed architecture to be used which is detailed later in the chapter (refer Section 4.2).

Document analysis for primary requirement analysis

Halstead (1977) and Wolverton (1974) state that requirements need to be systematically organised especially when they are large. This is because modelling and tracing requirements for large next generation systems adds to development costs while simultaneously increasing complexity.

LSDC's are enormous cyber-physical systems (CPS) with thousands of sensors feeding data to various automation and/or system controllers that can measure and control numerous variables of the data centre (Parolini, Tolia, Sinopoli & Krogh, 2010). This complexity of LSDC's makes analysing thermodynamic processes and modelling controllers a large and complex undertaking.

As per the guidelines proposed by Halstead (1977) abstracting LSDC's will reduce the complexity. The abstracted LSDC's is in the form of connected networks, and these networks are thermal, computational, physical, power etc. Figure 4.1 represents the various connected networks used to abstract LSDC's for simplistic modelling.

Focussing only on the thermal network while neglecting the other networks makes analysing and modelling the LSDC easier. Figure 4.1 illustrates the abstracted model to represent the connected network of LSDC's. The layers represent the different networks of the data centre such as network, computational etc., while nodes represent the different components in the layer. For example, in the thermal network, the nodes would be servers, heat exchangers etc. The links between the nodes represent the interactions between the components. For example, in the thermal network the links represent the interconnection of devices to illustrate the coolant flow in liquid cooled LSDC's.

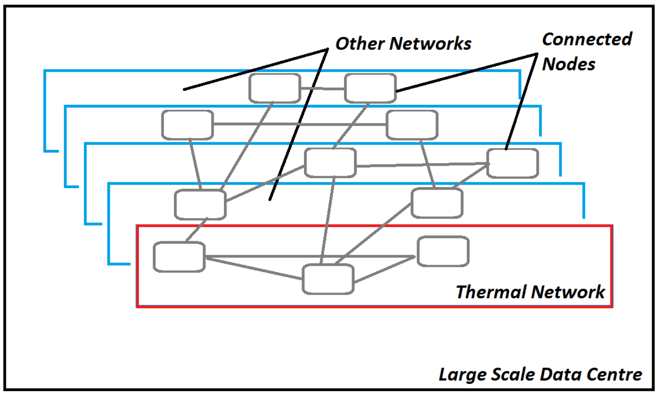

Figure 4.1: Abstracted model to represent the connected networks of LSDC's

LSDC's being cyber-physical systems have a clear divide amongst its devices and networks. This layered approach is based on the design challenges of CPS outlined by Lee (2008). Lee also states a key feature of CPS is the embedded nature and feedback loops which exists in a layer and the interaction between layers. These features are captured by the links. The networked nature of LSDC's shows us that there exists a collection of elements which interact and affect workings of each other. This implies other network layers such as the computational layer affects the thermal layer, which is accurate since the thermal output of a server rack is based on the IT load.

This model allows us to focus on control strategies for specific networks during development, although the scope for this thesis is only the thermal network. Unified Modelling Language (UML) is the standard used for modelling in model driven (D. C. Schmidt, 2006). The requirements of the DynaCool system are discussed in following section:

4.1.3 DynaCool system requirements

The need for proper requirements elicitation and the guidelines for formulating proper requirements was outlined in the literature review chapter (refer Section 2.6.4). As asserted by Adams (2015) requirements are classified into two paradigms; functional and non-functional.

Functional requirements

The DynaCool simulation software or in short DynaCool, helps data centre operation analysts/strategists to simulate various control strategies before implementing the same on physical LSDC's. The DynaCool software is designed to simulate different liquid cooling control strategies for any data centre topology configuration. This feature is required to test various intelligent control algorithms for implementation on Programmable Logic Controllers (PLC's) or intermediaries such as (Internet Of Things) IOT devices. IOT devices are internet connected hardware capable of performing various operations including controlling industrial systems.

The DynaCool system must provide necessary analytical data to users such that they make an informed decision before physical implementation. The minimum data an analyst needs is the total power consumption of cooling pumps, temperatures of server racks and the working coolant. These data metrics allow the user to calculate the optimal flow rate of the coolant for a given topology. Analysts can then optimise the cooling control strategies by reducing power consumption of coolant pumps, such that the temperatures of both the coolant and servers do not exceed the thermal regulatory guidelines set forth by ASHRAE (ASHRAE, 2011) for any given data centre class and usage.

A simple way to describe the functional requirements is to use a black-box diagram which shows the data requirements for the DynaCool system (T. Y. Chen & Poon, 2004). Figure 4.2 represents the black box diagram for the DynaCool system, which is used in association with non-functional requirements to facilitate the development and validation of the DynaCool system (Kumari, 2014).

Figure 4.2 illustrates the required data inputs and outputs for the DynaCool system, these data inputs and outputs are needed to achieve the operational requirement of the

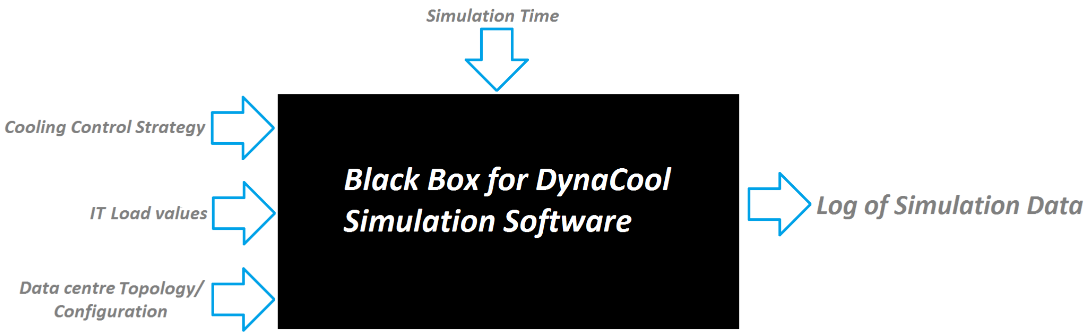

Figure 4.2: Black Box diagram showing the data requirements for the DynaCool system

system. The requirement is to provide simulated data showing the power consumption of the cooling pumps and the effect of cooling control strategies on the data centre topology. The data requirements in order to achieve the operational requirement are as follows,

Simulation time: The total duration for the simulation to take place, after which the DynaCool system stops and outputs the data for further analysis.

Cooling control strategy: The algorithmic implementation of the control logic being tested.

IT load values: The IT load values for each of the server racks in the given data centre topology, expressed in percentages where full and idle loads correspond to 100% and <15% respectively.

Data centre topology: The data centre model which shows the fluid dynamical relationship between the different data centre components.

Non-functional requirements

While non-functional requirements (NFR's) tend not to directly achieve system tasks, they are crucial for functional requirements. According to the International Standards Organization (ISO, 2010) and Ebert (1998), NFR's are *"Software requirements that describes what the software will do but not how the software will do it"*. According to

Chung, Nixon, Yu and Mylopoulos (2012) they are relative and interacting and evaluated subjectively rather than objectively, this is because they are difficult to quantify, are usually stated vaguely, applicable only to a particular system, and are global so they cannot be localized. Software performance objectives, design limitations, interface characteristics and quality are all examples of NFR's.

While the Institute for Electrical and Electronics Engineers (IEEE) have proposed guidelines for developing a System Requirements Specifications document (Board, 1998), they are not specific for NFR's. The IEEE Recommended Practice (Committee & Board, 1998) for NFR's is that they should consist of *"performance, functionality, design constraints, attributes and external interfaces"* and *these guidelines have been followed for formulating the NFR's of the DynaCool system.*

With these guidelines it is important to prioritise NFR's into levels which are an important measure for developers and management to categorize and implement requirements. The levels of priority are listed below,

1. Priority 1: This is the highest level, meaning this requirement is critical in the software development process and must be implemented into the product for many reasons like laws, policies, regulations etc.

2. Priority 2: This is the mid-level, which is a feature that is good to have and in most cases is implemented into the product as it poses an immediate benefit, but is sometimes left out due to constraints like budget, time etc.

3. Priority 3: This is the lowest level, these requirements are usually nice to have and is often implemented to enhance experiences of end users, for example a progress bar to notify users of the time remaining. These features pose no immediate benefit but rather make the final product more complete and professional.

4.1.4 Ordered Non-Functional Requirements (NFR's)

The NFR's listed below are based on the RE process outlined in previous sections from requirements gathering through document analysis, stakeholder analysis and suggestions from experts in the domain.

List of NFR's having priority 1	
Requirements	Page Ref.
The system should use CFD calculations for enhanced accuracy in simulating cooling of LSDC components	58
The system should be accurate and validated against published data	109
The system should be adaptable to different data centre configurations	58
The system should handle 'n' number of data centre infrastructure components such as cooling pumps, server racks, heat exchangers etc.	75
The system should be capable of getting input from an intelligent controller	77
The system should be clearly de-coupled to separate controller logic from data centre model i.e. it should be distributed	76
The system should be written in a language that can be ported to multiple platforms i.e. Platform independent for	54
The controller should have a latency of less than 500 milliseconds in reading the temperatures of the servers	109
The controller should be able to poll the data centre model for information twice as fast as the model is performing its calculations i.e. it should perform twice as fast as the model calculates The system should have negligible latency in the order of milliseconds	54 54
The system should function without the need for large computing resources	54

Table 4.1: Ordered priority 1 NFR's

NFR's are not complete without considering the constraints under which the system is developed and implemented, these are listed below.

List of NFR's having priority 2	
Requirements	Page Ref.
The system should be able to communicate without dependence i.e. it should be able to communicate remotely over the internet or via a local network (abstracted)	76 and 77

Table 4.2: Ordered priority 2 NFR's

List of NFR's having priority 3	
Requirements	Page Ref.
The system should provide feedback to the user, I.e. a view to display relevant information such as server rack temperature and flow rate of cooling pumps	90
The system should display prompts to users to load different experimental parameters to test the best control strategies	92

Table 4.3: Ordered priority 3 NFR's

Constraints

Constraints can be divided into 2 categories, Software and hardware with their own set of limitations. These constraints were formulated from exhaustive interviews and brainstorming sessions with experts from the data centre community. Experts who were instrumental in determining the set of constraints were from Microsoft Azure, a LSDC operator and cloud provider. Numerous other analysts with data centre infrastructure acumen, pitched in their ideas to help formulate a comprehensive set of constraints for the DynaCool system and the abstraction of cooling control logic.

Software Constraints

1. The system needs to operate across various platforms i.e. be platform independent.
2. The system needs to written in a language that can incorporate complex control logic in controllers, even have the ability to connect across interfaces to exploit cloud services like machine learning algorithms and complex statistical

approaches like regression.

3. The system needs to be portable i.e. not be installed, rather run 'out of the box'.

4. The system needs to be scalable.

Hardware Constraints

1. The system needs to operate with very little overhead and use a reasonable amount of computing resources.

2. The system needs to be easy to port onto PLC's or intermediaries.

3. The system needs to be backward compatible with older legacy computers/devices.

The software requirements formulated are accurate and proper but have not been formalized. This is because we do not require any processing of these requirements for automated model generation. We do not use them for any validation other than to communicate the software's intent to the developer, which in this instance is only us. Hence a gruelling formal process for NFR's would prove to be counter-productive and thus is excluded from the scope of this thesis.

Gomaa (1993) states that most literature on object oriented system designs (OOSD) omit important design issues when modelling real-time distributed applications. Unified Modelling Language (UML) is the standard for complex software design and is essential to blend OOSD with UML to develop distributed real time systems (Gomaa, 2001). Booch, Rumbaugh, Jacobson et al. (1999) explains, UML has standard notations for incorporating OOSD models which is applicable to the DynaCool and as such is the modelling language we shall use henceforth for system development.

4.1.5 Assumptions and Dependencies

While designing the system, certain assumptions had to be made since complete information of future or even current data centre's cannot be obtained due to access restrictions. Jaeger, Lin and Grimes (2008) explains this is because the information is sensitive and proprietary to data centre operators. One reason they state which resonates with the ideas expressed in this thesis is that the physical infrastructure is optimised to provide a high quality of service and is therefore sensitive information. The dependency of the software is due to the limited scope of knowledge which adheres to the previously identified constraints. C# and the .Net framework adheres to these constraints and is why it was chosen as the programming language and framework of choice for development.

Assumptions

1. The software assumes that pumps are connected in parallel using a 'T' junction, which sums the flow rates of individual pumps and series connections have no effect on net velocity of flow rate, this is based on the information published by Fernandez, Pyzdrowski, Schiller, Smith et al. (2002) on the operation of coolant pumps.

2. The software assumes basic computer literacy on the part of the user.

3. The software also assumes sufficient computing resources are available (Very little is needed and is met by all of computing hardware since 2000, first introduction of .net)

Dependencies

Operation of the software is dependent on the .NET framework released by Microsoft in the late 2000's.

4.2 Model Driven Engineering (MDE)

MDE allows for iterative development of models and controllers, which is done to refine the same. This section aims to detail the implementation for the DynaCool system design. Sinha (2013) shown MDE and model based design's prowess over other methods of designing complex coordinated systems. Their study on coordinated traffic controllers draw the same parallel to DynaCool processes which highlight the strengths of MDE. These strengths become evident in designing embedded systems simulation software especially with the use of Model-view-controller architectures.

4.2.1 Architecture selection

Due to the DynaCool system being modular and distributed in design a suitable architecture needs to be implemented. The IEC 61499 is one such architecture which is highly modular and efficient in its implementation (Dubinin & Vyatkin, 2007). The IEC 61499 uses function blocks as its primary method of providing object oriented design through encapsulation. IEC 61499 is well known for industrial programmable logic controllers which hosts the logic for liquid cooling systems.

However, this architecture lacks several crucial features as per the requirements identified, it does not possess methods of creating dynamic data centre configuration models as it relies heavily on function blocks (Dubinin & Vyatkin, 2007). The architecture also lacks the ability to add advanced control logic and instead relies mostly on custom algorithmic commands to implement logic, hindering management and deployment of more advanced control strategies in future LSDC's.

These are only a few of the drawbacks in using the IEC 61499 architecture even though it is recommended for embedded systems. After analysing numerous architectures such as Client-server, Service oriented, component-based etc. the most suitable architecture emerged, Model-View-Controller which satisfies all of our requirements and is the reason we have selected it for use in DynaCool. The IEC 61499 is applicable and as such some of its principles such as event-driven execution and feedback loops are used for the DynaCool system. The DynaCool system is therefore designed to be a modular software for distributed control systems through component re-use and encapsulation.

4.2.2 Model-View-Controller (MVC) implementation in DynaCool

The Model

The model in MVC is such that it is indifferent to the "outside world", meaning it is decoupled and only interacts with data, control etc. through interfaces thus when considering the model, it is important to consider its 2 parts, application and domain, both of these together constitute the complete model.

Domain Model: This is the kind of model that designers and analysts think of as a 'Model', this consists of objects which reinforce and embody essence of the problem (Deacon, 2009). The domain model consists of classes crucial to the core of the implemented solution, which is why it is highly structured and sometimes referred to as the Blue Book (Golberg & Robson, 1980). It contains mission-critical logic which compliments the Application model. In the DynaCool system design, this model incorporates most of the thermodynamic processes and calculations that occur in the data centre model simulation.

Application Model: The application model inherits most mechanisms from the domain model, it is mainly concerned with interfacing and supporting the domain model and is why it is sometimes referred to as the Application Co-ordinator (Deacon, 2009). There is a clear divide between the two models; the application model acts to facilitate the functions of domain model and holds objects rather than classes, again complimenting the domain model. This separation allows for robustness and flexibility, even though the application model interfaces, do not have the information about the views and controllers (K. Brown, 1995). In the DynaCool system design, this model incorporates the input output data operations. This includes writing simulation data to files, reading the data sent by the controller and processing the data while feeding it to the domain model to complete a simulation cycle.

UML activity diagrams for LSDC plant or facility modelling

Activity diagrams, according to the study published by Dumas and Ter Hofstede (2001) are intended to model workflows, both organizational and computational processes. They fail to capture some useful situations which suggest improvements are necessary. Since UML currently only supports activity diagrams, we shall use them to model our systems along with the other diagrams like Use Case, illustrated in Figure 4.4, to ensure a more accurate and complete model is implemented.

The UML activity diagram for the DynaCool system is illustrated in Figure 4.3. The DynaCool system is abstracted into 3 processes, known as Config, Model and Controller. The config process is used to ascertain the simulation parameters from the user and after successfully validating the parameters it automatically generates the data centre topological configuration and calls the model process to start the simulation. The main parameters the user needs to enter is the total simulation time, cooling control strategy also known as algorithm preference and the size of the IT load values to be written. The Cooling control strategy is one of three, PVFR, VFR and static; these have been alluded in the literature review chapter (refer Section 2.5).

The model process simulates the thermodynamic interactions between the various components in the data centre topology. The model calculates in real time the thermodynamic interactions and writes the simulated data to files which are later read by the controller. The model also reads and processes the data sent by the controller. Controller modelling is detailed in the following section.

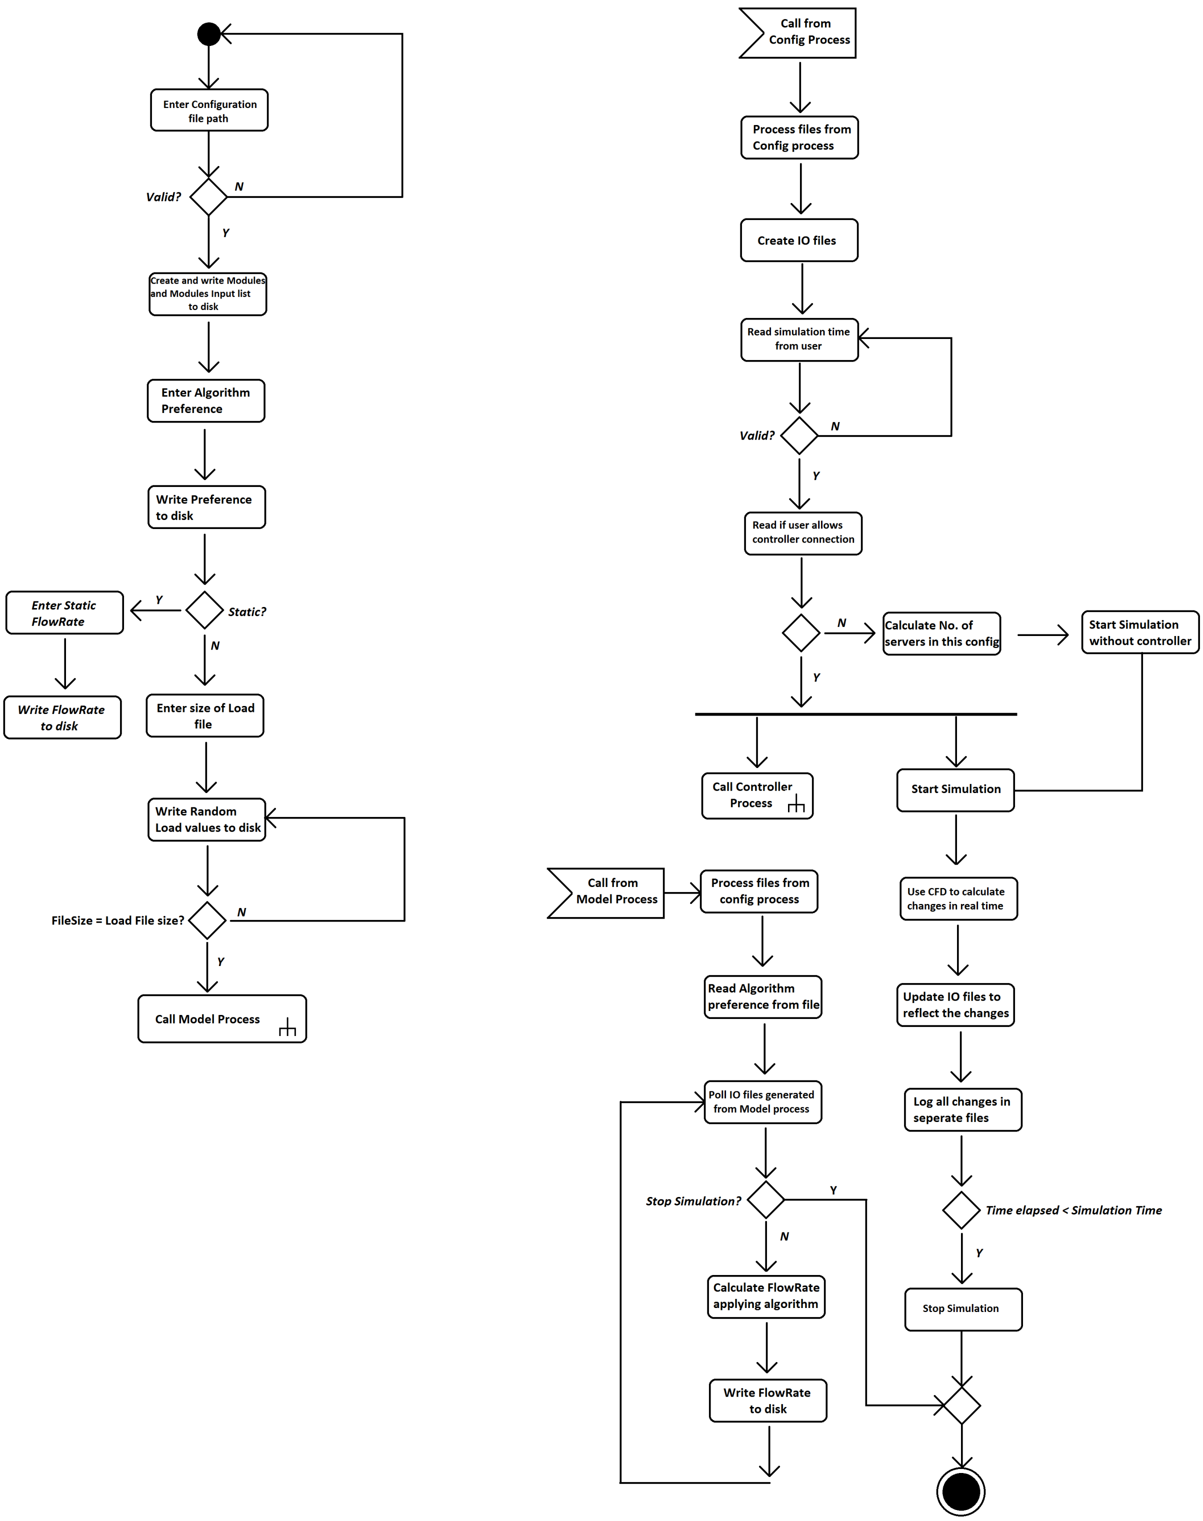

Figure 4.3: UML Activity diagram for DynaCool system

The Controller

This is an object which can perform manipulations. Consider this as a manager, managing user input or model output and input. Controllers have the most knowledge about the underlying hardware, and are usually reusable. Controllers can be replicated and distributed in their implementation (Patterns, n.d.; Deacon, 2009).

The controller in the DynaCool system is abstracted to be distributed and concurrent with any other controllers that may be trying to interface with the model. This is done because controllers may fail and adding redundancy measures is critical in mission critical systems, in which category liquid cooling systems are present.

The controller process has embedded cooling control algorithmic logic. The controller calculates the optimal coolant flow rate based on the cooling control strategy and logic. The data is processed from the model data and the optimal flow rate is sent to the model process to actuate the coolant pumps. Modelling cooling control strategies are illustrated in the sections below (refer Section 4.2.4).

The View

This is intuitive, representing the display. It is the only point of contact for the user in most cases. Examples of views can be one of the three, Graphical User Interface (GUI), Command Line Interface (CLI) and Application Programming Interface (API). This is the interface to which input and outputs are implemented against. There maybe one or more views in an application (Deacon, 2009).

In the DynaCool system design, the view is merely to get the input parameters for the simulation from the user and to give relevant simulation information. There are two views, one for the model process and another for the controller process which are running throughout the simulation. The initial view, the configuration process as mentioned before is only to get the simulation parameters from the user after which it is closed.

Considering all of the benefits mentioned the MVC architecture is ideal for our application, however the features of IEC 61499 could prove beneficial including as function blocks and clocks, and as such they have been incorporated into the final version (V2.1) of the DynaCool system. This is because sensors in cyber-physical systems such as data centres provide data in discrete intervals or periods which reduces the polling burden on systems when fully implemented on physical hardware. This improves performance significantly, and is one of NFR's identified. Using function blocks also increases distributiveness. Distributing processing loads helps mediate risk and reduces complexity. This is because redundancy and risk can be spread across multiple devices if need be, and these features are supported by MVC.

4.2.3 Use Case diagram for system modelling

Use Case diagrams are used to present a graphical high level overview of the functionality of a system, they are mainly used to analyse the possible interactions of actors (users, customers, managers etc.) with the software. Use Case diagrams are also used to represent the relationships of various components present in the system (Gomaa, 2001).

The Use Case diagram for the DynaCool system is illustrated below in Figure 4.4, the diagram shows the direct interaction of the user between the config and model processes. The user enters simulation parameters to generate accurate simulation data which is later used by data centre operation analysts.

4.2.4 Modelling the different cooling control strategies

PVFR

The novel Pulsed Variable Flow Rate (PVFR) is an intelligent cooling control strategy which incorporates our hypothesis that using pulses of power instead of continuous

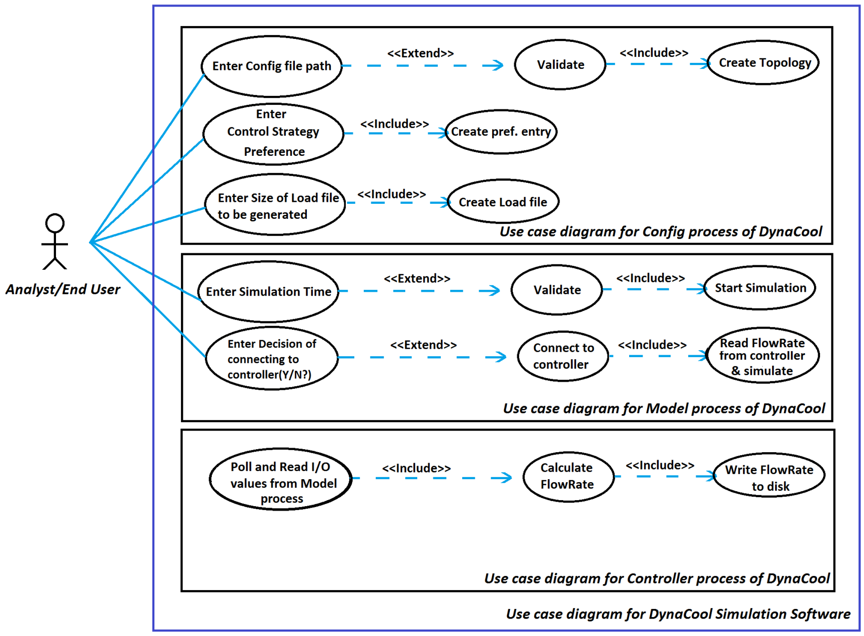

Figure 4.4: UML Use case diagram for the DynaCool system

power supply reduces the power consumption of the coolant pumps and increases efficiency of the data centre. The approach is to reduce power consumption by maximising efficiencies of all the components involved in liquid cooling systems.

Our premise is that coolant pumps operate at maximum efficiency at 80-90% of full load (Majidi, 2004), and operating them at such a load all of the time leads to over-cooling and under-cooling of servers. We also intend to exploit the advancements in thermal variance tolerance of microprocessors (Intel, 1998). This theory of using pulses rather than continuous supply of power has been tested and proven in other areas of electronics such as lighting. Choi et al. (1994) have shown to increase battery life in flashlights by optimal pulsing of light emitting diodes.

We intend to apply these principles for energy reduction in single phase closed loop liquid cooling systems used in some LSDC's. The main advantage of implementing these systems is that it allows for modularity and rapid scalability for LSDC operators. Another key feature of a PVFR control strategy allows LSDC operators to dynamically adjust the coolant flow rate based on the demand placed on servers, which allows for zoning and economical infrastructure utilisation.

VFR

Variable Flow Rate (VFR) is an intelligent cooling control strategy which follows the Google's AI approach of continuously varying the coolant flow in real time to meet current cooling demands of servers (Gao, 2014). This approach has been explained in detail in the literature review chapter (refer Section 2.5).

Static

The static cooling control strategy employs a non-intelligent approach, where the coolant flow is calibrated and set to a fixed coolant flow rate. This approach is also discussed in the literature review chapter (refer Section 2.5).

4.3 Summary

In Section 4.1, existing solutions for were explored for the problem of cooling in LSDC's. Discussion of stakeholder analysis and requirements elicitation processes for the design of DynaCool system. The functional and non-functional requirements for the DynaCool system were outlined along with hardware and software constraints. The non-functional requirements were ordered and prioritised in three levels to allow for simplifying the development process. Assumptions and dependencies needed for DynaCool development were also detailed.

Section 4.2 discussed the model driven engineering process used for system design, the UML activity and use case diagrams were illustrated to provide comprehensive system design details for further development of the DynaCool system. The section also briefly discussed the control strategies involved and our novel approach for designing the PVFR cooling control strategy.

Chapter 5

Building DynaCool

This chapter focuses on specifying the technicalities involved in accurately modelling and building the DynaCool simulation software. DynaCool is designed to provide accurate simulation data to test different cooling control strategies for LSDC operation analysts and strategists.

Sections 5.1 and 5.2 detail the development of the LSDC model and the controller respectively. The controller incorporates the desired cooling control strategies as discussed previously in section 4.2.4 of this thesis. The organisation of the different sections of this chapter is based in part on the MDE development guidelines proposed by Smiałek and Nowakowski (2015). The iterative development of the LSDC model is necessary for the controller to operate, which is why the development of the LSDC model takes precedence over the controller. Section 5.3, summarises the chapter by explaining the development of the DynaCool system.

5.1 LSDC model design

5.1.1 Mathematical Model Selection

Although CFD analysis has been used throughout the industry, validity and accuracy of these systems remain largely unverified. This is especially true when predictions are made on large scale data centres (Shrivastava et al., 2006; Almoli et al., 2012). That being said, proper mathematical model selection plays a pivotal role in determining the final accuracies of the system. This was evidenced when Almoli et al. (2012) opted to use a Reynolds Averaged Navier-Strokes (RANS) model to simulate air flow for cooling in server racks. This decision was based partly on a study by Jinkyun et al. (2009). The model neglected to take into account the heat transfer rates and thermal states of the server rack and fluid, leading to erroneous results.

We have opted for a universal approach using thermodynamic equations coupled with Reynolds number calculations. These calculations help us to formulate accurate fluid dynamical equations which are used to determine if the fluid flow is turbulent or laminar. This approach is similar to using Large Eddy Simulation (LES) models. As Gousseau, Blocken and Van Heijst (2011) state, LES models provide greater accuracy for stream-wise turbulent mass transport over RANS model to predict convective fluxes.

The argument that LES models are more accurate than RANS models is contested by Jinkyun et al. (2009). They assert RANS models provide better accuracy in some instances when compared with LES models. This statement was tested by Almoli et al. (2012); who concluded, improper application of any model will prove to be erroneous. This suggests LES models are indeed more accurate, but only in some specific instances RANS models may prove to be more accurate. Gousseau et al. (2011) further analysed the accuracy of LES models in turbulent flows and concluded that they are more accurate in turbulent flows when compared with RANS models. In our use case, the fluid flows are turbulent in nature (refer Section 5.1.2) and applying LES models in such an instance offers significantly better approximations as ascertained by Gousseau et al. (2011).

5.1.2 Building the Computational Fluid Dynamic (CFD) model of the thermal network

Figure 5.1 shows the schematic of the fluid flow dynamics between the various components in a 1U rack-mount server. The schematic is based on the modular design proposed and patented by Hastings et al. (2002), which is used by all server racks in LSDC's.

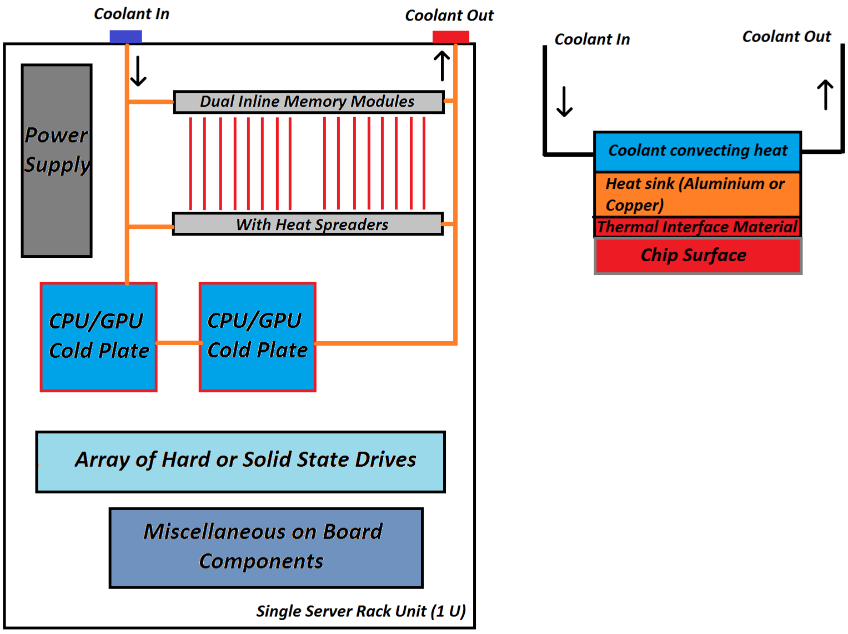

Figure 5.1: Schematic of the fluid dynamic model for a 1U rack-mount server

A heat sink is a small amount of material used to facilitate heat transfer from a chip surface to a working fluid such as water or air. To accurately model heat transfer between the coolant and CPU or GPU, we need to first account for thermal resistances. The total thermal resistance is the sum of resistances of the thermal interface material and the aluminium. Copper is a better thermal conductor when compared to aluminium, even so we have chosen aluminium/ This is because aluminium forms an oxide barrier which makes it non-corrosive, making it ideal for use in cooling applications. This non-corrosive property is especially critical in cooling applications where the coolant is highly oxidative, for example water.

Equation 5.1 represents the total thermal resistance in the system.

$$R_{Total} = R_{TIM} + R_{AL} \tag{5.1}$$

Thermal conductivity is an important property to consider, as it has a significant impact in cooling applications (Eckert & Drake Jr, 1987). The equation for thermal conductivity is given in Equation 5.2.

$$ThermalConductivity = \frac{Thickness}{R_{Total} \times Area} \tag{5.2}$$

Solving Equations 5.1 and 5.2,

R_{Total} = 0.2 + 0.02 = 0.22 k/W and Thermal Conductivity = 0.0101 / (0.22 × 0.0001) = 0.000022 W/(m.k)

Modelling the heat transfer

Please refer the literature review chapter of this thesis (refer Section 2.2.4) to understand the different modes of heat transfers that can occur in a thermodynamic system. The mode of heat transfer will be convection for the DynaCool system as we are using a flowing coolant to cool server rack components. Diffusivity of the fluid is a characteristic of the fluid used to determine the rate of heat transfer. To calculate the dominance of diffusivity of fluid flow, to determine the type of convection for our LSDC system model we need to understand the type of convection occurring in the system. This is needed as it will affect the heat transfer characteristics between the components of the system and is calculated using the Reynolds and Prandtl numbers. These numbers are representative of the characteristic of the fluid flow and are dimensionless (Eckert &

Drake Jr, 1987). Convection is a function of these numbers as illustrated in Equation 5.3, and the relationship between them is given by a Nusselt number equation illustrated in Equation 5.4 (Ellison, 1984). Nusselt number is the ratio of convective to conductive heat transfer and given by the Dittus-Boelter equation for turbulent flows as illustrated in the Equation 5.5.

$$Convection = f(Re, Pr) \tag{5.3}$$

$$Nu = \subset \quad Re^{x} \times Pr^{n} \tag{5.4}$$

Where, Re is the Reynolds number, Pr is the Prandtl number, x is the Linear Dimension (m) and n is 0.4 when heating of the fluid and 0.3 when cooling of the fluid.

Nusselt numbers are needed in our model as there is a conductive heat transfer from the server components to the heat sinks and a convective heat transfer from the heat sinks to the coolant.

$$Nu_D = \quad 0.023 \times Re_D^{4/5} \times Pr^{n} \tag{5.5}$$

Where, D is the inner diameter of the pipe (m).

The equations to calculate Reynolds and Prandtl numbers are given in Equations 5.6, 5.7, 5.8 and 5.9.

$$Re = V_\infty \times L/V \;\; or \;\; Re = [(\rho L^3)(V^2/L)]/\mu(V/L)L^2 \tag{5.6}$$

Simplifying Equation 5.6 we get,

$$Re = \rho \nu L/\mu = V \times L/\nu \tag{5.7}$$

Where, V_∞ is the Maximum velocity of fluid (m/s), ρ is the Density of fluid (kg/m^3), μ is the Dynamic viscosity of fluid (kg/m.s) and ν is the Kinematic viscosity (m^2/s)

As the coolant flows through a pipe in the DynaCool LSDC model, the equation to get the Reynolds number is given by Bernouli's equation (Ellison, 1984) (refer Equation 5.8),

$$Re = Pr \times D_H/\mu \quad = v \times D_H/\nu = \quad Q \times D_H/v \times A \tag{5.8}$$

Where, D_H is the Hydraulic diameter of the pipe (m), Q is the Volumetric flow rate (m^3/s), v is the mean velocity of fluid (m/s) and A is the cross sectional area of the pipe (m^2)

$$Pr = v/\alpha = \mu \times C_p/K \tag{5.9}$$

Where, μ = Dynamic viscosity (kg/m.s), K = Thermal conductivity (W/m.k), C_p = Specific heat (J/kg.k)

The coolant Properties of the fluid (water @ 20 $°C$) according to data published by Poole, Sciortino, Essmann and Stanley (1992) are as follows,

α is 3000, ρ is 1 kg/m^3, μ is 1.002, C_p is 4.184 kJ/kg.k and K is 0.6 W/m.k @ 20 ° C

Solving Equations 5.7, 5.8 and 5.9, Re = 1 × 1 / 1 × 0.0001 = 10,000

Pr = 1.002 × 4.184 / 0.6 = 6.987 @ 20 ° C

These numbers indicate that we have turbulent flow and that momentum diffusivity dominates. This is because Reynold's number > 4000 is classified as turbulent and Prandtl number > 6 indicates momentum diffusivity dominates (Ellison, 1984). Therefore, based on the nature of fluid flow we can determine that the type and mode of heat transfer is one of *'Forced convection'*.

The generalized equation for convection can be calculated by using Newton's law of cooling for forced convection (Bergman & Incropera, 2011), which is illustrated in Equation 5.10.

$$\frac{\partial Q}{\partial t} = h \times A \times \Delta T(t) \tag{5.10}$$

Simplifying Equation 5.10 yields,

$$Q = h \times A \times \Delta T \qquad\qquad when \quad t = 1 second \tag{5.11}$$

Where, Q is the thermal energy (J) , A is the Surface area (m^2), h is the heat transfer co-efficient ((W/m^2)-k) and Δ T is T - (T_{env}) (k or $^\circ$ C) i.e. the difference in temperatures between the heat flux and convection fluid (water, air etc.)

Equation 5.11 represents the real time working calculations that need to be performed in order to calculate temperature decreases.

Modelling the Heat Exchanger

The DynaCool system is a closed loop system, and as such the heat absorbed by the coolant from the server racks need to be dissipated efficiently to ensure safe operation of the system. This heat dissipation can be accomplished by using a heat exchanger. The function of a heat exchanger is to transfer heat from one working fluid to another. This means it is a device which facilitates heat transfer from a higher temperature fluid (coolant) to a lower one (waste cooling fluid). Note that working fluids may be gases or fluids (Kakac et al., 2012).

To determine the proper heat exchangers to suit our application, we need to calculate the rate of heat transfer (Bergman & Incropera, 2011) for a given type of heat exchanger. The equation to calculate it is illustrated in Equation 5.12.

$$Q = U \times A \times \Delta T_m \tag{5.12}$$

Where, Q is the heat transfer rate, U is the overall heat transfer co-efficient, A is the heat exchanger area and ΔT_m is the mean temperature difference between the fluids.

For any heat exchanger, there is a minimum separation distance between the two working fluids, this can be calculated by the equations (refer Equations 5.13 to 5.20) given below,

Heat transfer rate from a hot fluid to the wall area Δz is given by Newton's law of cooling as:

$$q = h_{c1} \times A_s[T_H - T_1] \quad such\ that,\ \ T_H - T_1 = q/h_{c1}.A_s \tag{5.13}$$

Where, h_{c1} is the heat transfer co-efficient (W/($m^2.k$)), A_s is the heat transfer area in m^2 and T_H and T_1 are the temperatures of the working fluids.

The temperature drop across the wall can be found by solving the conduction equation (refer Equation 5.14 and 5.15),

$$d/dx(K \times dT/dx) = 0 \tag{5.14}$$

subject to boundary conditions,

$$K \times (dT/dx)|_0 = q/A_s \tag{5.15}$$

Where, d/dx and dT/ dx = differentiation functions and K = Thermal conductivity of the body and q = flow capacity in m^3/h.

Integrating Equation 5.14 and applying boundary condition yields,

$$\int_0^{x'} d/dx(K \times dT/dx)dx = 0 \tag{5.16}$$

$$K \times dT/dx' - K \times dT/dx|_0 = 0 \tag{5.17}$$

$$K \times dT/dx' + q/A_s = 0 \tag{5.18}$$

Integrating Equation 5.18 over the wall thickness, we obtain the temperature drop across the wall which is given below,

$$\int_0^L dT/dx'dx' + \int_0^L q/K \times A_s dx' = 0 \tag{5.19}$$

$$T_1 - T_2 = q \times L/K \times A_s \tag{5.20}$$

Equation 5.20 represents the temperature drop of coolant when it passes through the heat exchangers, which is used as a real time calculation in the DynaCool system.

Modelling Coolant Pumps

A coolant pump is essential in liquid cooling systems to provide adequate coolant flow to cool the LSDC IT components. The power consumed by a coolant pump is given in Equation 5.21 (Rezania et al., 2012).

$$P = q \times \rho \times g \times h/(3.6 \times 10^6) \tag{5.21}$$

Where, P is the power rating of the pump in kW, q is the flow capacity in m^3/h, ρ is the density of fluid in kg/m^3, g is 9.81 m/s^2 and h is the differential head (m).

Pumps often are stated in terms of mass flow rate, therefore we need to convert mass flow rate into volumetric flow rate to get the rating of our pump being used. The equation for converting mass flow rate into volumetric flow rate is as below,

$$Vol.FR = mass.FR \times \rho \tag{5.22}$$

Where, ρ is the density of the fluid, Vol.FR is the volumetric flow rate in m^3/s and mass.FR is the mass flow rate in kg/s.

The equation for converting volumetric flow rate into m^3/hr is as below (refer Equation 5.23). This calculation is required to determine the power rating of the pumps needed to provide a flow rate of 10 kg/s for a given LSDC topology.

$$FlowRate = \frac{Volumetricflowrate \times 3600}{1000} \tag{5.23}$$

Solving Equation 5.21 by assuming an optimal data centre height of 3.65 meters or 12 feet we get,

$$P = \frac{3600 \times 1000 \times 9.81 \times 3.65}{3600000} = 35.81kW$$

35.81 kW or 48.02 Horse-Power (HP) is the required size of the coolant pumps to provide a mass flow rate of 10 kg/s, which is the pump size used in the DynaCool LSDC model. This is based on the LSDC model published by Patel and Shah (2005).

Load modelling

There are two different types of loads in a LSDC, IT load and thermal load. The IT load is measured in percentages and represents the total demand of computing resources placed on a server. 100% IT load represents full load and <15% represents the idling of servers. Thermal load is a function of IT load, an idling server still expels up to 66% of the heat produced if fully loaded (G. Chen et al., 2008). we have modelled this characteristic into the DynaCool models by calculating idle heat outputs and full load heat outputs of server racks in real-time based on the IT load.

5.1.3 Modelling the LSDC topology

We have discussed the LSDC topology and its effects on the thermal profiles of servers in the literature review chapter (refer Section 2.2.4). In this section we will focus on modelling the topology from a thermodynamic perspective. There are two main paradigms that are essential to discuss namely, the cooling infrastructure of LSDC's and the topology of the server racks and servers.

The purpose of the cooling infrastructure is to remove heat expelled from the server racks for maintaining safe operating parameters. In this instance, the DynaCool LSDC model is a closed loop single phase liquid cooling system, and as such modelling of the two cooling loops are essential. These loops are illustrated in the Figure 5.2. They are a loop of coolant circulation to cool the server racks and another to cool the circulating fluid.

These loops are affected by the topological configuration of the server racks, because of the thermodynamic interactions that occur for different configurations. According to De Groot and Mazur (2013) in a non-equilibrium system, interactions between the heat source and convective fluids can vary with configurations. In our instance, an increase in server rack power density might have different thermal characteristics over increases in the number of server racks. This will remain the case even if the total heat load placed on the cooling system remains the same. Server racks typically host 42, 1U servers (Hastings et al., 2002). These server racks are usually daisy chained in single phase liquid cooling systems (Sickinger et al., 2014).

The cooling infrastructure for this application provides immediate coolant cooling, which can only be achieved by using a liquid-liquid counter flow shell heat exchanger (Kakac et al., 2012; Primo, 2012). The system uses a cold fluid which is circulated opposite to the direction of convection coolant flow. The heat absorbed by the cold fluid is exactly equal to the heat lost by the coolant. The efficiency of these shell heat exchange systems are high but consume more power when compared to free-air cooling systems which are currently being used by other LSDC operators (Google, 2016a). This is why the DynaCool system and the cooling control strategies employed will be tested against these variations in topologies through experimentation.

The DynaCool LSDC model resembles that of Asetek's RackCDU system (Sickinger et al., 2014). The RackCDU system has already been implemented and tested. The DynaCool model has an on-chip single phase set up with a thermal interface material to facilitate the heat transfer process.

Schematically the DynaCool LSDC system model is illustrated below in Figure 5.2.

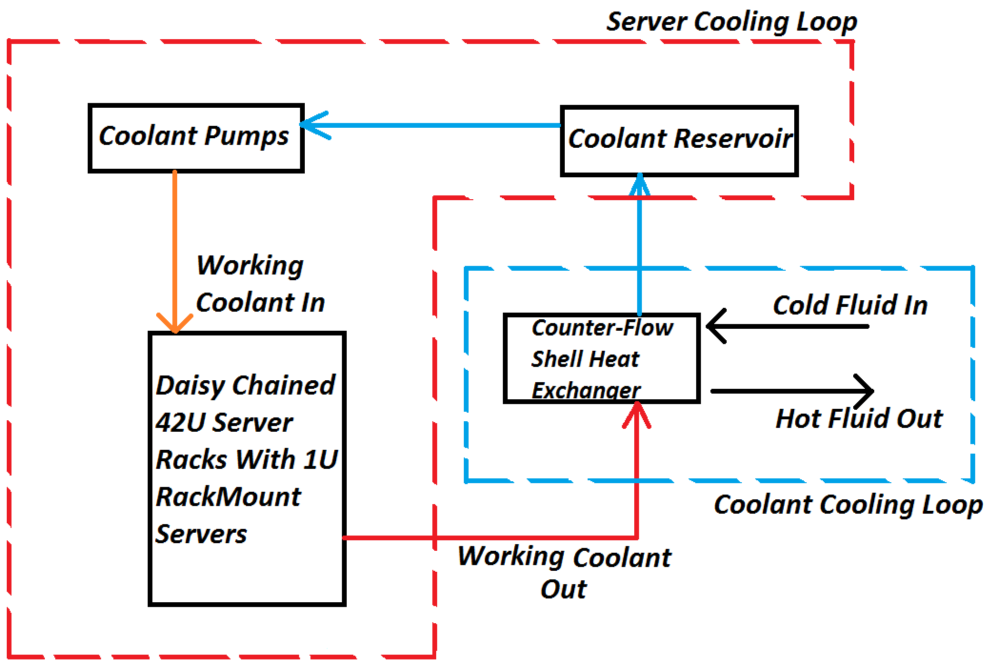

Figure 5.2: Schematic of the DynaCool LSDC system model

5.2 LSDC controller modelling

In order to deliver a controlled coolant flow, pump control systems use Pulse Width Modulation (PWM) for varying flow rates (D. J. Allen & Lasecki, 2001). Flow rates in coolant pumps are measured using mass flow rates (Majidi, 2004). This means a certain mass or amount of coolant is being pumped per second of operation measured in kg/s. PWM is a method of delivering impulses of power to a system, the working of PWM as a control strategy for pumps is illustrated in Figure 5.3.

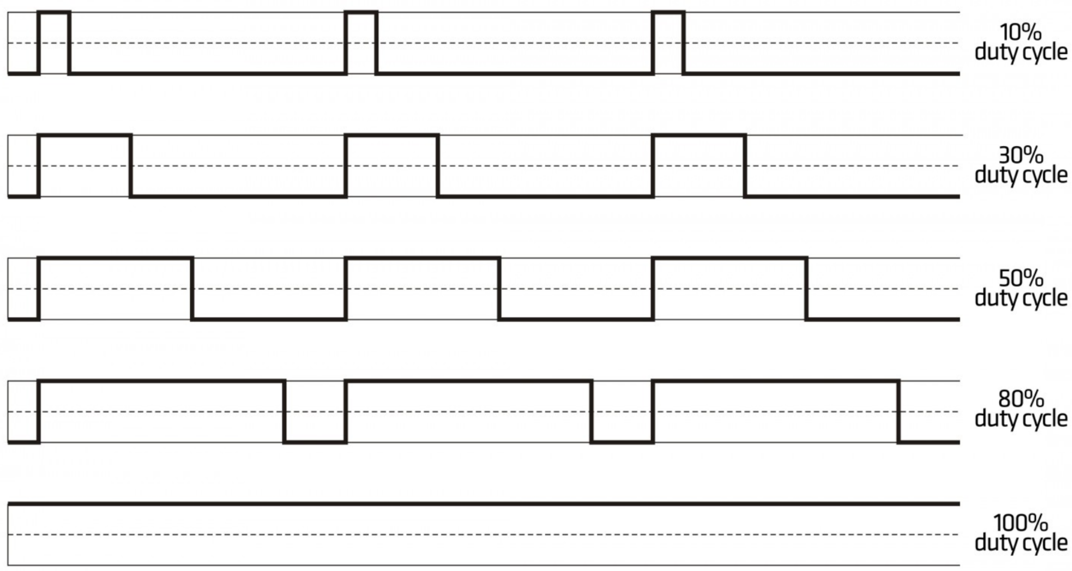

Figure 5.3: Working of PWM as a control strategy for pumps
adapted from (Gobor, 2016)

Our hypothesis is that, using a pulsed variable flow rate (PVFR) control strategy offers enhanced efficiency when compared to a continuously varying flow rate (VFR) control strategy. The efficiency is based in terms of power consumption of the coolant pumps to maintain safe operating temperatures of the servers. The temperature range of these servers are maintained according to ASHRAE (2011) guidelines. To test the hypothesis, we will simulate the power consumptions of three different control strategies, PVFR, VFR and static flow rate.

5.2.1 Version 1 controller

The first version of the controller is the algorithmic implementation of the novel PVFR cooling control strategy. This version was based on a temperature based control strategy where the coolant flow begins when the temperature of servers reaches a higher threshold (T2) and shuts off when it reaches a lower threshold (T1). The temperatures T2 and T1 are set according to the recommended guidelines published by ASHRAE (2011).

This version during initial testing showed promising results for reducing the power consumption of pumps but failed to meet the performance requirements of the systems. The control strategy worked well for a small number of server racks, but when extended to meet a real world use case of a module of ten server racks the processing delays added up to make the implementation infeasible.

The PVFR version 1 algorithmic implementation of the cooling control strategy is illustrated below in Figure 5.4:

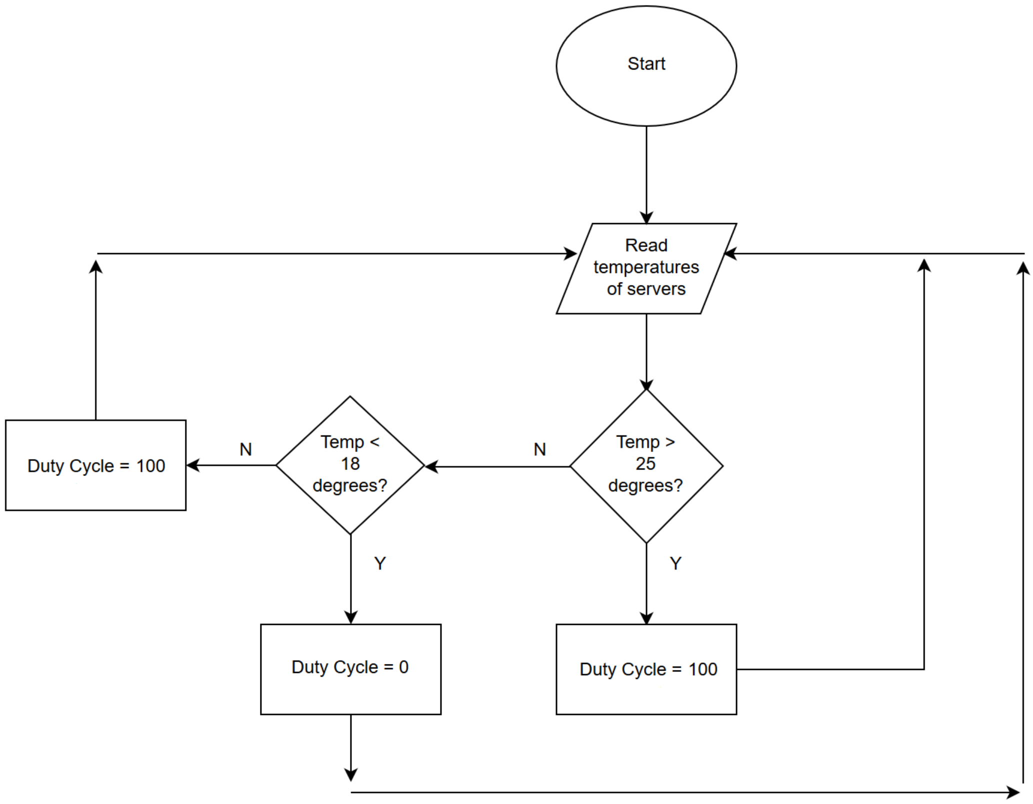

Figure 5.4: Flow chart illustrating the working of PVFR v1 cooling control strategy

The PVFR version 1 controller reads the temperatures of servers and checks if it exceeds 25 $°C$ the maximum temperature (T2). If the temperature exceeds T2 then the duty cycle is set to 100 and is only set to 0 if the temperature reaches below 18 $^{\circ}C$ the minimum temperature (T1). T2 is set 2 $^{\circ}C$ below the allowed temperatures guidelines by ASHRAE because it ensures safe operating parameters and counteracts against the latency of the coolant pumps to provide coolant flow. Also, as seen in Figure 5.4 there is no 'end'. This is because the controller is designed not to end its execution which is critical to prevent fires in emergency situations. The safety measure set in place is to set the duty cycle to maximum which should help prevent fires.

5.2.2 Version 2 controller

The second version of the PVFR cooling control strategy aimed to eliminate the processing burden that occurred in version 1. The way this was achieved was to pulse the coolant flow till the next temperature read cycle i.e. every second. This ensures that the computation resources required and the delays in reading the temperature sensors of all the servers are significantly reduced.

The PVFR version 2 algorithmic implementation of the cooling control strategy is illustrated in Figure 5.5:

The PVFR version 2 controller reads the temperatures of servers and performs the same checks as version 1 controller, but actuates the coolant pumps every second rather than every 0.1 seconds, as was the case in version 1 controller. In this version only temperature T2 is significant and since the actuations are only one second long, the problem of over-cooling and under-cooling are eliminated. This is because the controller can dynamically adjust the coolant flow every second to either pulse the coolant or not to pulse the coolant. This eliminates under-cooling and over-cooling respectively.

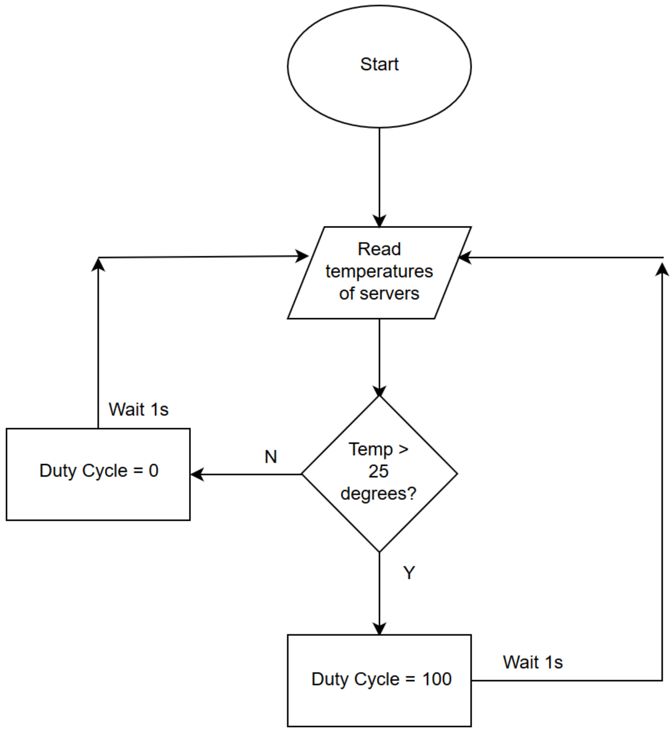

Figure 5.5: Flow chart illustrating the working of PVFR v2 cooling control strategy

5.2.3 DynaCool validation

Validation of simulation software is one of the most important measures that can be taken to ensure the results obtained are valid. The two ways of testing the DynaCool system are done through a facility cooling failure test or commonly referred to as a ride through test and a steady state test. During these tests the cooling systems are shut down to measure the temperature rise of servers and a constant flow of coolant to measure the temperature drop of servers respectively.

Since power densities of current data centres are 30 kW (R. Schmidt et al., 2005) we performed a ride through test with a single server rack of 30 kW. The results we got from the DynaCool system was a 1.2 $^{\circ}C$ rise in temperature per second of operation per server rack, the results mirror testing analysis done in **refer Section 6.1.1**. A steady state test where the server racks were provided with sufficient cooling is also analysed in **refer Section 6.1.1**. These results are in close agreements with the data published by (Moss, 2011; Intel, 1998; Madhavi, 2014) in their physical data centres. Both the tests used server racks with power densities of 30 kW. After the ride through test, the steady state static cooling tests and their measurements were found to closely agree with the data published by Subramanian (2008).

Furthermore, we analysed the heat exchangers and found that the heat dissipation or temperature drop of the coolant are in close agreement with the results laid out by Primo (2012). The temperature of the working coolant fluid when the mass flow rate remained constant at 6 Kg/s, and the inlet coolant fluid was 10 $°C$.

The heat sink temperatures of the DynaCool system were in close agreement with the results published by Intel (1998). The results published by Intel were tested on physical devices at a much smaller scale, hence we used to scale power densities to compare results under the same testing criteria.

Finally, the pump calculations were analysed and compared with the results published by Mathew Milnes (n.d.) and Kang et al. (2007), we found that the results of the pump and the thermal characteristics of the system closely agreed with the published data. Also the latency between the duty cycle input to the desired coolant flow can be calibrated using the polling time variable in the controller logic, enhancing the accuracy of the system to meet real world use cases. Currently, this latency is 500 milliseconds, based on the calculation outlined by Fernandez et al. (2002).

The thermal characteristics of the DynaCool system was also independently modelled using AutoDesk's simulation software, a commercial CFD modelling and analysis tool. The limited properties that could be analysed like temperature rise and heat exchange characteristics were analysed and compared with commercial data from (Moss, 2011; Kang et al., 2007; Subramanian, 2008; Mathew Milnes, n.d.).

These results were also cross references with the data published by Autodesk (2013), and the data published by Intel (Musilli & Ellison, 2012) they were close agreement and hence we can safely conclude that the *DynaCool system is accurate and valid*.

Accuracy of Simulation Software

A cooling system is used to reduce the temperature of a given substance which is why cooling systems are dependent on the thermal loads placed on them. This was believed to imply a linear relationship between heat loads and the amount of cooling required (Chandler, 1987). However, this assumption leads to erroneous results, especially when considering server racks where high thermal dissipation occurs in a confined space. Our statement is supported by Patel et al. (2002), when they conducted a study for high thermal dissipating servers, they concluded larger cooling systems require complex thermal management systems and in practice, they do not perform linearly and such an assumption would lead to erroneous results. This is because a linear assumption would lead to inaccurate data generation by the simulation software and any control strategies made using this data will lead to system failure.

Dai et al. (2014) draw a similar conclusion based on their testing with air cooling and hybrid cooling systems. A linear assumption is inaccurate because such models assume an equilibrium model of thermodynamics. After exhaustive studies done by De Groot and Mazur (2013), a non-equilibrium model is shown to be a more accurate representation of the real world. There are two paradigms for the study of thermodynamics, namely Classical and Modern thermodynamics. The modern theory of thermodynamics focuses more on molecular fluid flow and their interactions. Liquid cooling is heavily concerned with fluid flows and the thermal interactions between the fluid flow and heat source. This is why opting for a modern theory of thermodynamics will provide the best probability of designing accurate models for use in the DynaCool system.

In the book, ‘Non-equilibrium thermodynamics’ authored by De Groot and Mazur (2013), we see a much clearer picture of fluid dynamics and the contiguous interaction of molecules with thermal energy. In non-equilibrium thermodynamics, the prevailing notion is that thermodynamic processes are irreversible and is based on the balance equation for entropy. Modelling the DynaCool system at a molecular level of detail is necessary because no other simulation software exists which simulate liquid cooling. DynaCool will set the precedence for the level of accuracy in liquid cooling simulation software and hence would prove to be worth the effort in detailing to such a level of accuracy.

That being said, molecular interactions are notoriously hard to calculate or predict which is why CFD simulation software are approximations of probable interactions. This statement is supported by studies performed by Shrivastava et al. (2006) and Almoli et al. (2012). They found that simulation software offers close approximations and most probable interactions but are not 100% accurate and should not be taken as absolutes. This is why we have enforced certain measures and have custom-tailored the software for our particular use case application to enhance accuracy. The use case being liquid cooling in data centres. Even with these measures, DynaCool is not 100% accurate and should not be considered as such.

5.3 Summary

In Section 5.1, we modelled the thermal network of a LSDC using a CFD approach. Section 5.1 also discusses the effect of topologies on the thermal profiles of the LSDC components using mathematical formulae. Section 5.2 focused on modelling the controller and its algorithmic implementation of the cooling control logic. Section 5.2 also detailed the evolution of the novel PVFR cooling control strategy. It also discussed the validation of the system models to ensure accurate results can be obtained using the DynaCool system.

Chapter 6

Experimental Results

This chapter focuses on detailing the experimental setup and extracting any knowledge or insights from the data gathered using the DynaCool system. Data will be gathered by testing three different control strategies namely, Static Flow Rate, Variable Flow Rate (VFR) and Pulsed Variable Flow Rate (PVFR). The data obtained from the system is the energy consumption of all coolant pumps for a given data centre topology and a given cooling control strategy.

Section 6.1 details the experimental setup and the data centre topologies used to gather simulation data from the DynaCool system. Section 6.2 focuses on the execution of simulations and illustration of the gathered data. Section 6.2 also details the method of execution of the simulations for the different experimental setups. Section 6.3 explains the specifics of the statistical tests used for data analysis on the gathered raw data. Section 6.4 describes the observations and insights recovered from the data analysis. It also discusses the limitations of the novel PVFR cooling control strategy and Section 6.5 provides a summary of the chapter.

6.1 Experimental setup

The thermal network of a data centre is greatly dependent on the topological configuration of its components. We have discussed the significance of the topological configuration on the thermal profile of the data centre while building the DynaCool system in previous chapters of this thesis (refer Section 2.4.1). The validity of the topological models being used were also explained in previous chapters (refer Section 5.2.3), which are based on real world models proposed by Sickinger et al. (2014) and Marcinichen et al. (2014).

We will conduct three different experiments to analyse the power consumption relationships between the three control strategies namely, PVFR, VFR and Static. Experiments are performed keeping all other variables like initial boundary conditions, data centre topologies etc. constant, so as to minimise erroneous results and biases. A daisy chained topology for the server racks is used. This means the output of one server racks coolant is the input for the next server rack.

Daisy chaining is done because it accurately represents the real world models identified by Madhavi (2014), Marcinichen et al. (2014), Autodesk (2013), Musilli and Ellison (2012) and Kang et al. (2007). The initial boundary conditions are the simulation parameters used to execute the simulations which are as follows,

- The server rack power densities vary for different experiments and hence are expressed in Sections 6.1.1, 6.1.2 and 6.1.3.
- The remaining boundary conditions remain constant for all experiments and are as follows.
- The maximum flow rate of the coolant pump is 10 kg/s.
- The capacity of the coolant reservoir is 200 litres and the room temperature is 15 $^{\circ}C$.

To ensure the experimental data is accurate and valid, randomised IT loads for each iteration are used to eliminate any experimental biases that may have been introduced knowingly or unknowingly. Randomisation and minimisation of the number of measuring variables (Treasure & MacRae, 1998) guarantees the results are not obtained by chance, allowing us to determine if the PVFR cooling control strategy is truly more efficient over the VFR control strategy.

Treasure and MacRae (1998) assert that minimization should be the 'platinum standard', by this they mean that, practitioners should minimize the number of measuring variables to determine the effects of an intervention over another. This is because unforeseen randomness can creep into experiments. The measuring variable in this instance is only one; power consumption of coolant pumps.

Randomisations and larger sample sizes reduce data variations and biases. These are standardised measures needed to remove any ambiguity and follow rigorous scientific procedures. This method when compared to observational studies is a pragmatic approach (Fagiolini et al., 2016). Using this approach with proper statistical analysis techniques is accepted by the scientific cohort as a valid method to determine the efficacy i.e. effectiveness of any interventions.

6.1.1 Setup used to determine the efficacy

The topology of the data centre used for this experiment is visually represented in Figure 6.1. We run seven iterations of the simulation using the DynaCool system for all the three control strategies. Each iteration used random IT load values placed on the servers while keeping the topology and other variables constant. The simulation will run for a total of 10 minutes for each iteration, and then an extended run time simulation will be performed for a total of 60 minutes.

The extended run time simulation is performed to establish the validity of the cooling control strategies under a continued time demand. It also proves to safeguard the control strategies legitimacy in a real world scenario, where they are implemented to perform continuously.

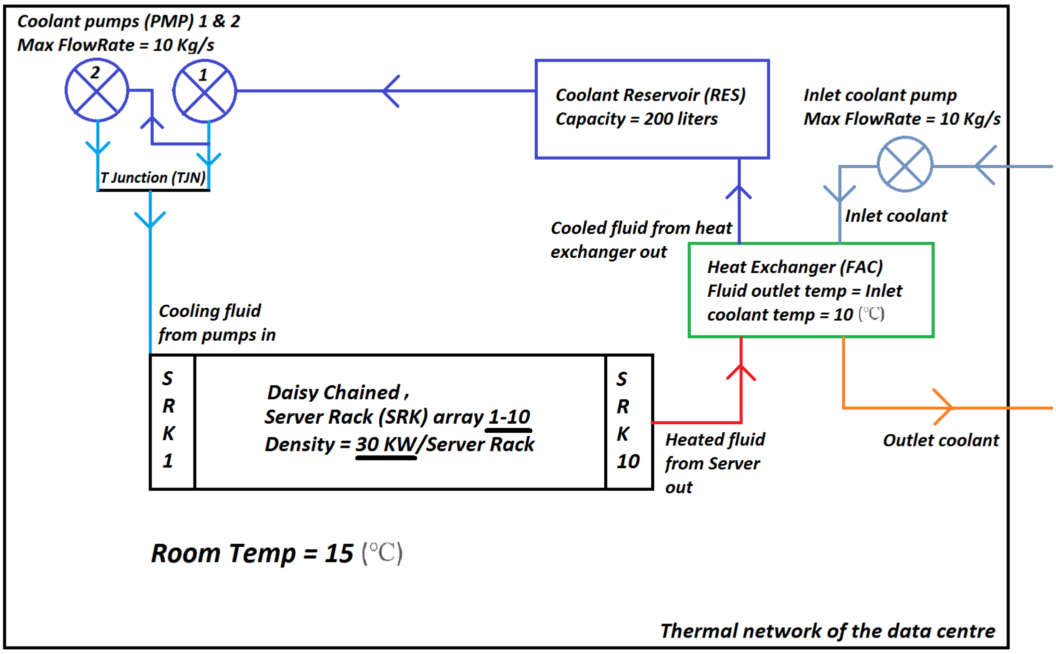

Figure 6.1: Schematic of data centre topology used while determining the optimal static mass flow rate and proving the efficacy of the PVFR cooling control strategy

6.1.2 Setup used while increasing the number of server racks

The topology of the data centre used for this setup is visually represented in Figure 6.2. The mass flow rate of the static cooling control strategy was set to 9 kg/s. This is because any flow rate lower than that threshold proved to be ineffective in cooling the server racks due to fact that increase in heat load requires an increased coolant flow. The simulation will run for a total of 10 minutes.

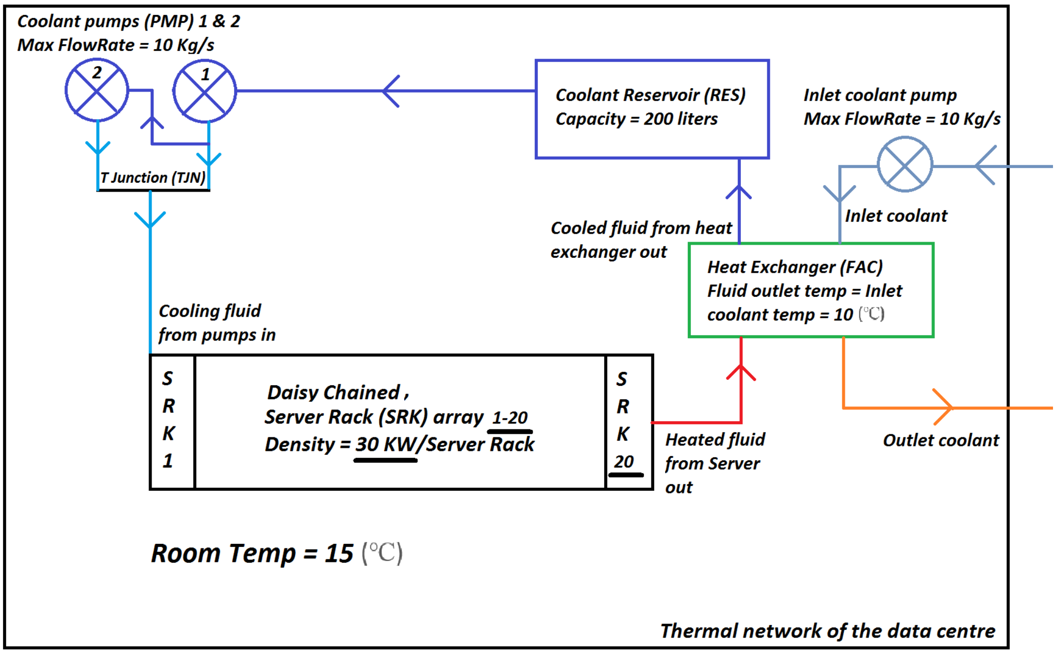

Figure 6.2: Schematic of data centre topology used while determining the effect of increasing the number of server racks has on efficiency

6.1.3 Setup used while increasing the density of server racks

Topology of the data centre used for this test is represented visually in Figure 6.3. The mass flow rate of the static cooling control strategy was set to 9 kg/s. This is because any flow rate lower than that threshold proved to be ineffective in cooling the server racks due to fact that increase in heat load requires an increased coolant flow. The simulation will run for a total of 10 minutes.

6.2 Execution and data gathering

This chapter details the simulation parameters used and the data gathered visually from all the three different cooling control strategy for the various experiments executed. Section 6.2.1 discusses the selection of an optimal static cooling control flow rate for the data centre topology illustrated in Figure 6.1. A static flow rate is calibrated

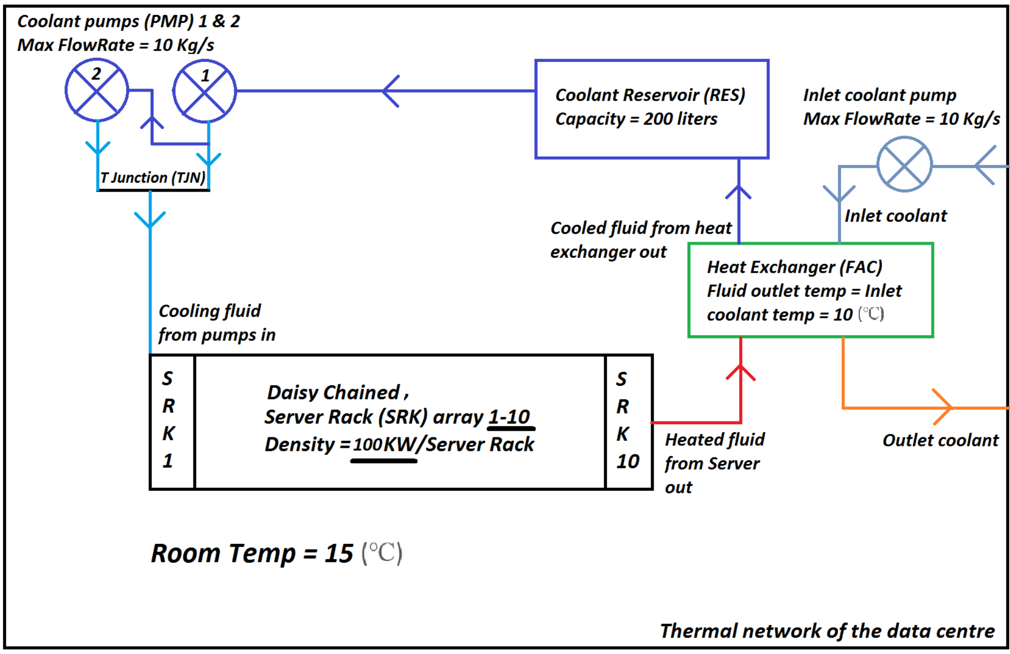

Figure 6.3: Schematic of data centre topology used while determining the effect of increasing the server racks density has on efficiency

to ensure a steady state simulation is obtained for validation as discussed in Section 5.2.3. Section 6.2.2 discusses execution of the simulation to determine the efficacy or effectiveness of the three different cooling control strategies. Section 6.2.3 and 6.2.4 details the simulation of increased server rack density and increased number of server racks respectively. Analysis of gathered data from the experimentation is performed in Section 6.3 of this chapter.

6.2.1 Determining the optimal static flow rate

A static flow rate control strategy employs a steady, unvarying flow of coolant in order to cool the servers. Finding the optimum flow rate for a given topology is usually done through calibration by trial and error.

We can narrow down the domain space by using the VFR and PVFR cooling control strategies. The cooling control strategies yield various informational data like maximum and minimum mass flow rates, which can be used to make an informed selection of a static mass flow rate. Figures 6.4 and 6.5 illustrates graphically the results of selecting various different static flow rates in terms of the power consumed by the coolant pumps and the server rack temperatures respectively.

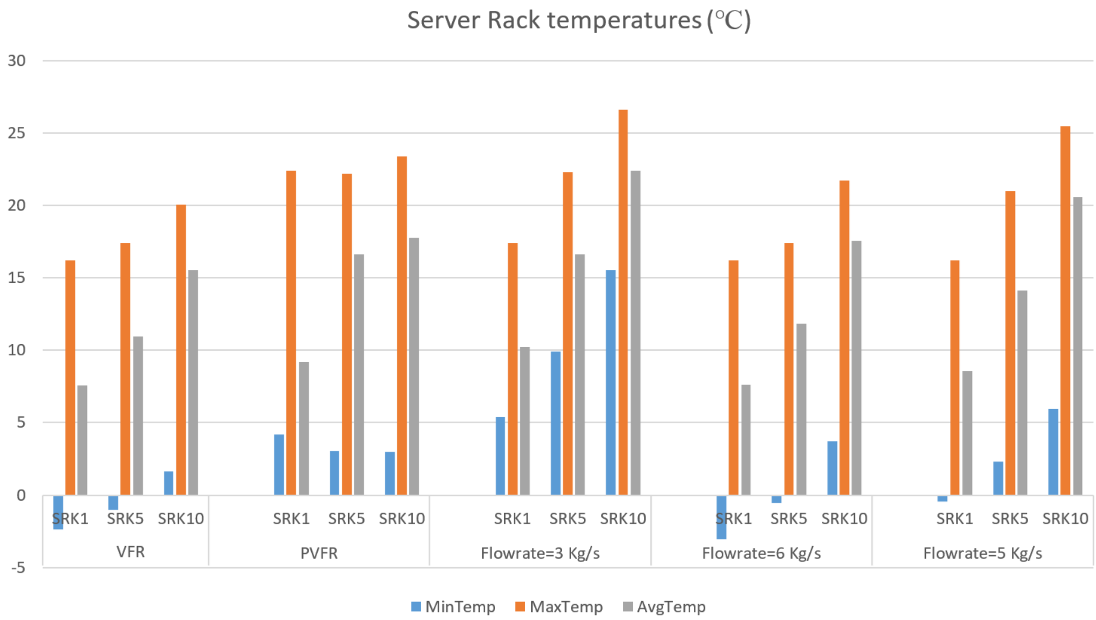

Figure 6.4: Server rack temperatures measured while finding the optimum static flow rate in degrees Celsius

Analysis

ASHRAE's (2011) recommended temperature guidelines for Class A (1-4) data centres is a maximum temperature of 27 $\circ C$. Referring to Figure 6.4,we can see that the maximum temperatures for the server racks 1,5 and 10 is below that threshold when a static flow is set to either 3 kg/s or 5 kg/s. Although a static flow of 5 kg/s should suffice in this instance, it exceed our predetermined buffer threshold set at 2 degrees below 27 $^\circ C$ i.e. at 25 $^\circ C$, for this reason we set the *optimal static mass flow rate at 6 kg/s for this particular configuration in all further experiments*.

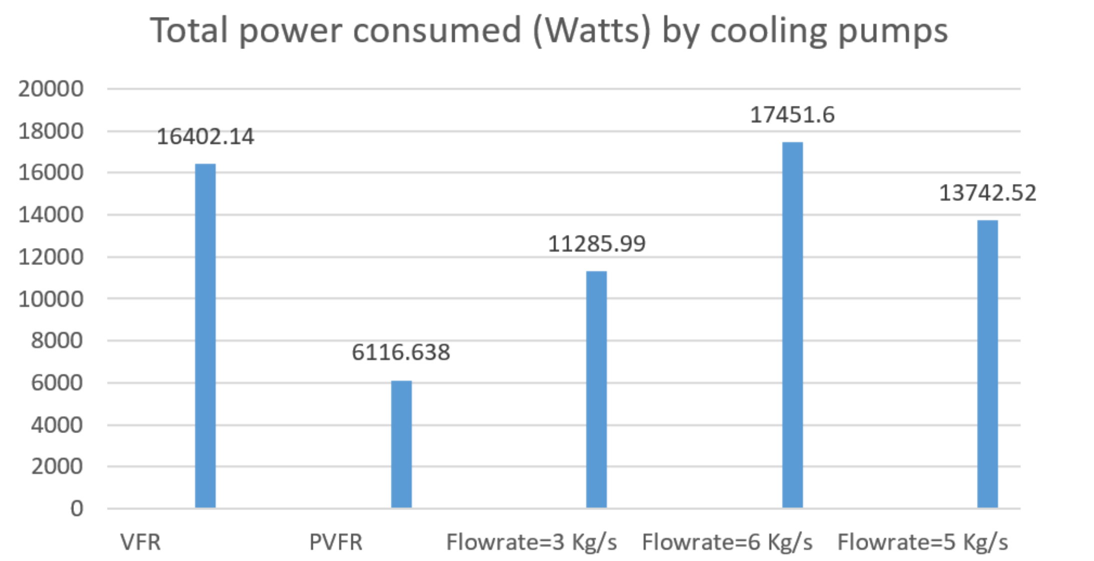

Figure 6.5: Total power consumed by coolant pumps measured while finding the optimum static flow rate in Watts

6.2.2 Experimental data gathered while determining the efficacy of a cooling control strategy

The raw data along with the DynaCool simulation software can be downloaded from the internet using the URL/link in the Appendix (refer Appendix A.1). Data was gathered using the DynaCool simulation software by performing the simulations based on the test setup criteria discussed in the previous sections (refer Section 6.1.1). Figures 6.6 and 6.7 illustrates the mean total power consumed by the coolant pumps for each of the three control strategies for a simulation time of 10 and 60 minutes respectively. Refer Figures A.1 and A.3 in Appendix A, for the corresponding server rack temperatures for the three different control strategies refer Figure A.2 in Appendix A, which illustrate the total power consumed by the coolant pumps in each of the seven iterations.

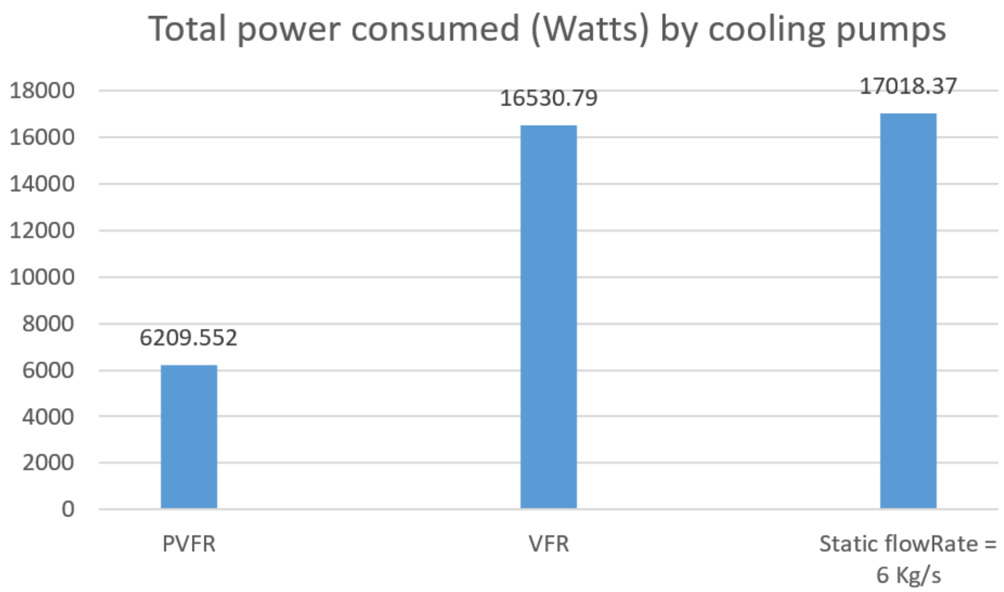

Figure 6.6: Mean Total power consumed by coolant pumps for the three control strategies while testing efficacy

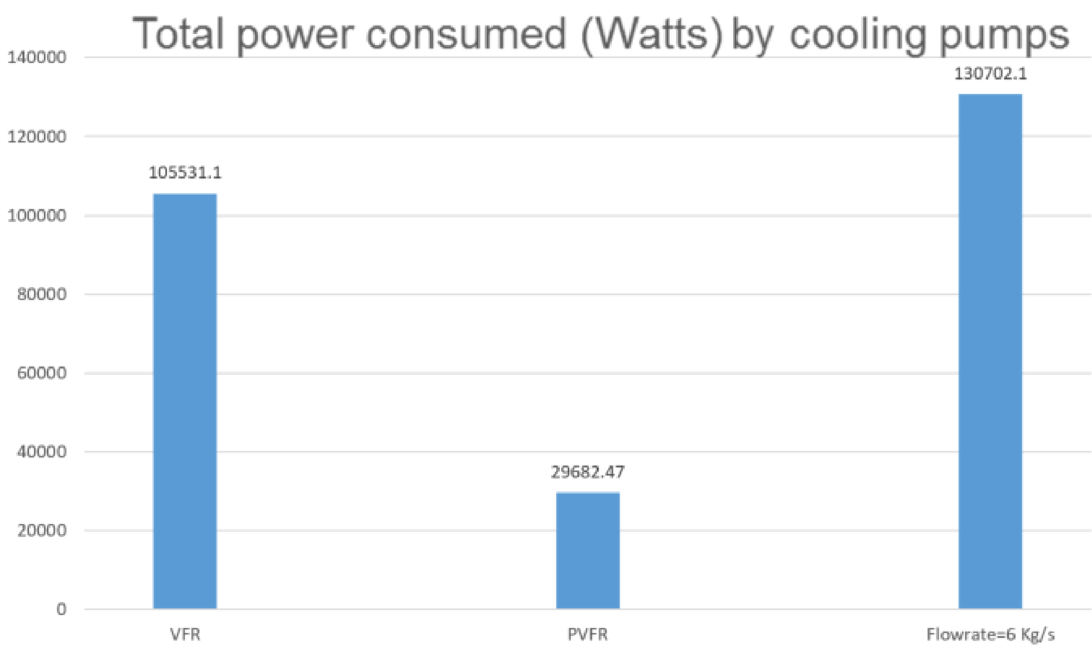

Figure 6.7: Total power consumed by coolant pumps for the three control strategies for a simulation time of 60 minutes

6.2.3 Experimental data gathered while increasing the server rack density

Data was gathered using the DynaCool simulation software by performing the simulations based on the test setup criteria discussed in the previous section (refer Section 6.1.2). Figures 6.8 and 6.9 illustrates the server rack temperatures for the three different control strategies and the mean total power consumed by the coolant pumps for each of the three control strategies respectively.

Increasing the server rack density has an effect on both temperatures shown in Figure 6.8 and power consumed shown in Figure 6.9. This is similar to the effect seen when simulating the increase in the number of server racks. As evident from Figure 6.11, the temperatures of the server racks exceed the ASHRAE recommended threshold (ASHRAE, 2011) of 27 $^{\circ}C$. The ASHRAE threshold is exceeded only when the VFR control strategy is implemented. As evident from Figure 6.9, the power consumption of the PVFR control strategy over static and VFR control strategies is 38.2% and 34.26% respectively.

Efficiency of the PVFR control strategy in this configuration is further reduced when compared to the preceding simulation. This is expected since the total heat load of the system has increased dramatically to 100 kW/server rack × 10 server racks = 1000 kW or 1 MW compared to 300 kW for the previous simulation. The maximum pump capacity of 10 kg/s is sufficient enough to cool the current configuration but not any further additions or increase in server rack numbers or power densities. This means the *upper end of cooling for two mass flow rated pumps of 10 kg/s is 10 daisy chained server racks at a power density of 100 kW per server rack.*

Another noteworthy point is that for future LSDC's a VFR control strategy such as the one proposed by Google which uses machine learning cannot be implemented at larger server rack densities. This is evident from the temperature readings of server racks illustrated in Figure 6.8 which exceed the ASHRAE (ASHRAE, 2011) thermal guidelines. This is because the *VFR control strategy* cannot effectively vary the flow rate based on rapidly varying load conditions, which is indicative of a real world usage scenario. Google's machine learning algorithms can predict future loads based on trends of past loads and taking into account usage trends etc. which is not possible in a *PLC driven implementation of cyber physical systems, which is a key limitation.*

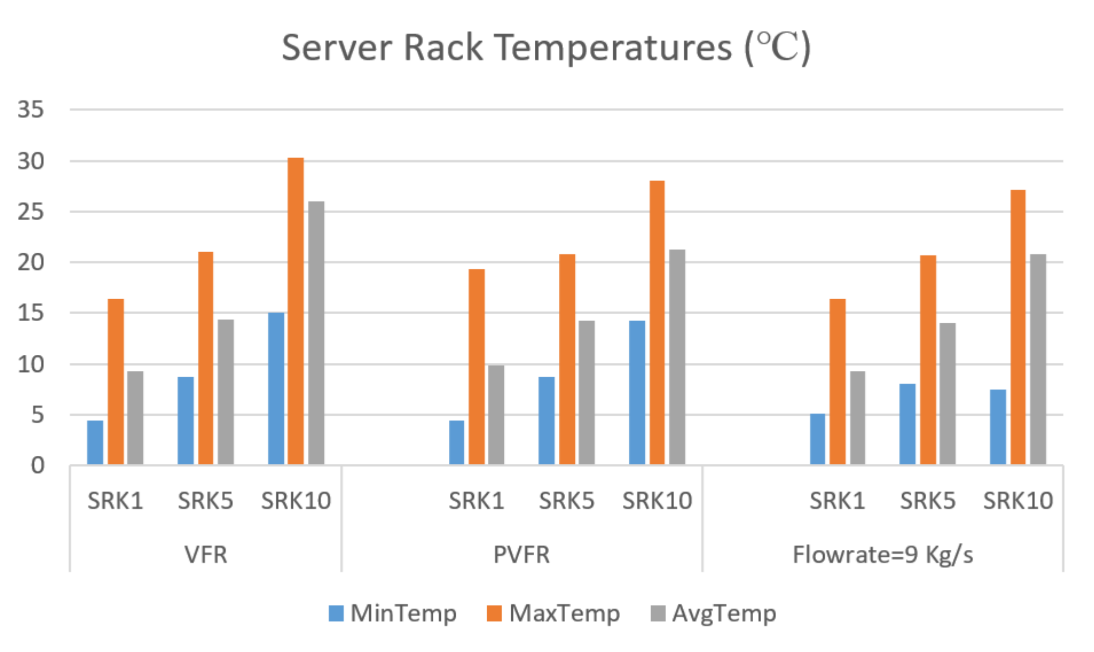

Figure 6.8: Server rack temperatures measured while simulating the increase in server racks density

6.2.4 Experimental data gathered while increasing the number of server racks

Data was gathered using the DynaCool simulation software by performing the simulations based on the test setup criteria discussed in the previous section (refer Section 6.1.3). Figures 6.10 and 6.11 illustrates the server rack temperatures for the three different control strategies and the mean total power consumed by the coolant pumps

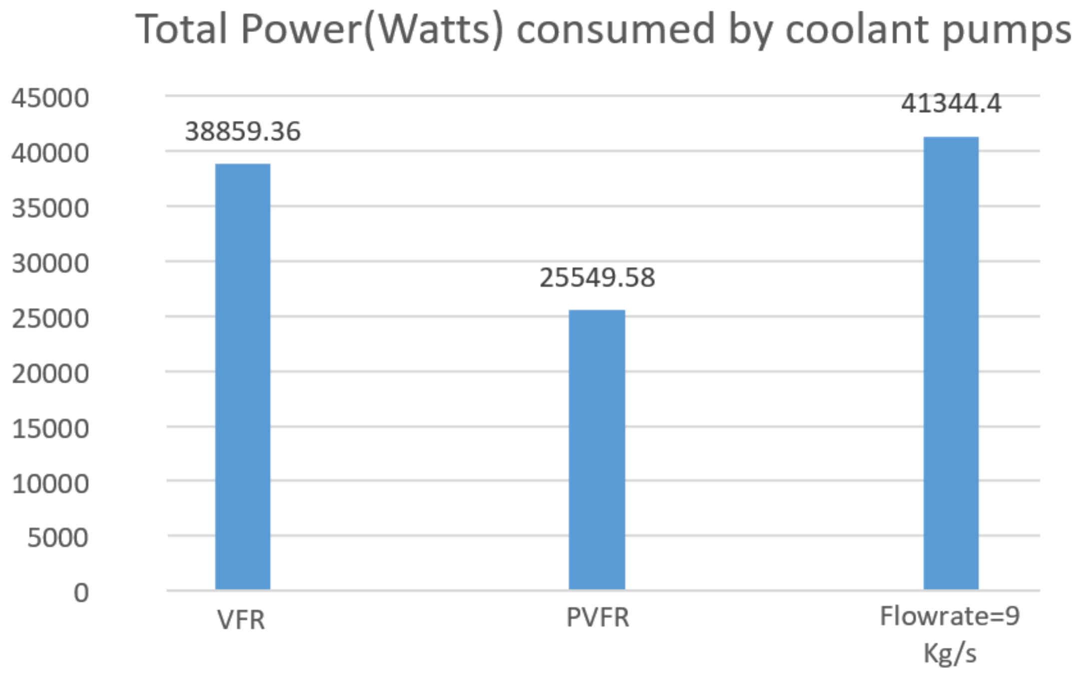

Figure 6.9: Total power consumed by each of the different control strategies measured while simulating the increase in server racks density

for each of the three control strategies.

As evident from the Figure 6.10, we can see maximum temperatures of all the three control strategies exceeding the buffer threshold of 25 $^{\circ}C$ but remaining below the 27 $^{\circ}C$ guideline set by ASHRAE (ASHRAE, 2011). We can see differences in the total power consumption of the PVFR to static and PVFR to VFR control strategies are 39.99% and 25.65 % respectively.

The efficiency of the PVFR control strategy in this configuration is reduced significantly when compared to the first simulation test. This means we can confidently deduce that there was an increase in total power consumed. This is expected since the total cooling power in this instance equates to 30 kW/Server rack × 20 Server racks = 600 kW. Compared to the first simulation heat load of 30 kW × 10 Server racks = 300 kW, this is effectively twice the heat load. Another interesting point to note is that, the maximum pump capacity of 10 kg/s is sufficient enough to cool the current configuration but not any further additions in server racks. This means the *upper end of cooling for two mass flow rated pumps of 10 kg/s is 20 daisy chained server racks at a power density of 30 kW per server rack.*

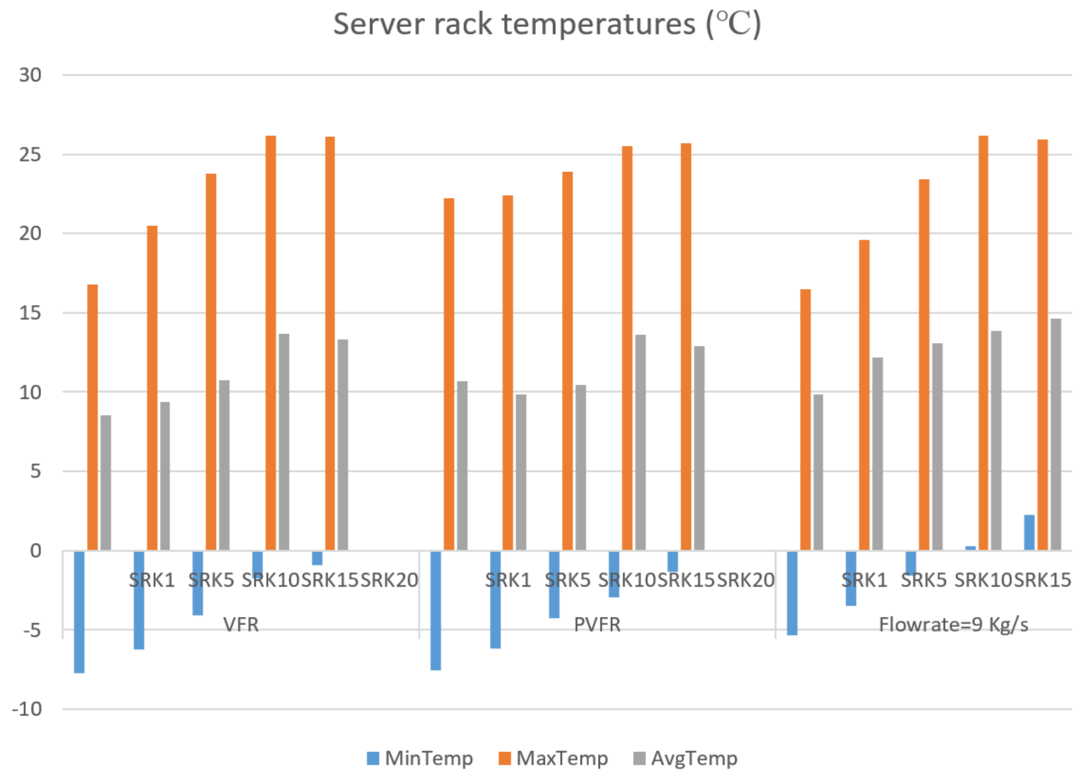

Figure 6.10: Server rack temperatures measured while simulating the increase in number of server racks

6.3 Data analysis

Data analysis is done using t-tests to eliminate and rule out any biases and ambiguity that may have been introduced during experimentation by random chance (Rice, 1989). This is why selecting a proper statistical analysis method is important and the process we have undertaken is detailed in the following section.

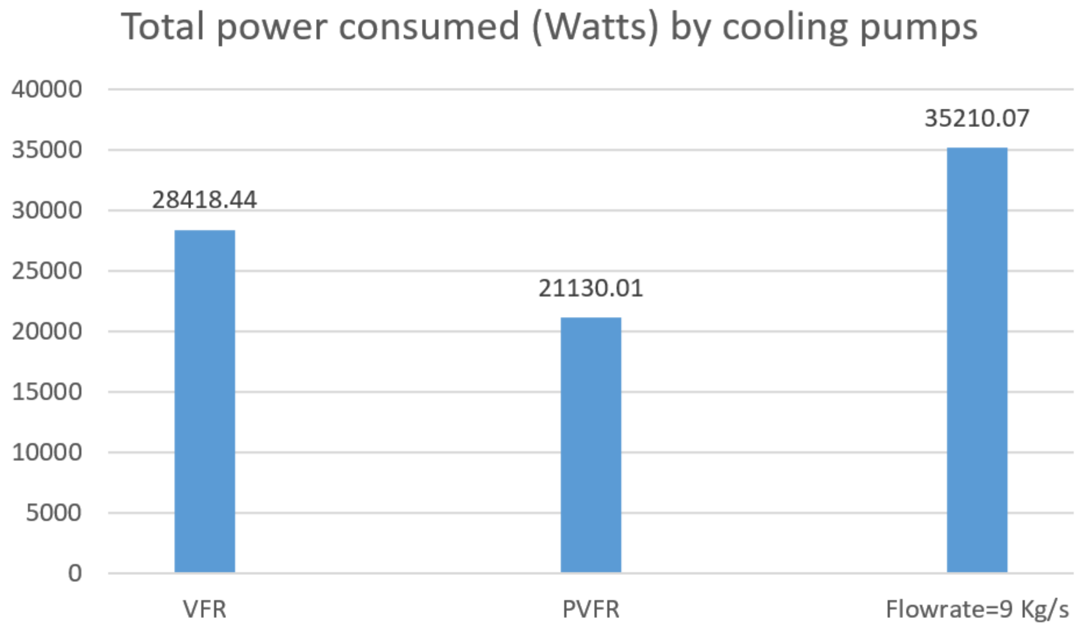

Figure 6.11: Total power consumed by each of the different control strategies measured while simulating the increase in number of server racks

6.3.1 Choosing a statistical analysis method

The alternatives to t-tests are chi-square, G and ANOVA tests (Mason et al., 2003). These are mainly used in evolutionary studies to identify regression, apply a curve fit and to calculate probability measures (Box et al., 1978). Chi-square, G and ANOVA tests employ complex statistical analyses which add to the complexity of a study. Chi-square, G and ANOVA tests are primarily utilised to predict future data based on past trends (Rice, 1989). In our study we do not use prediction models, which is why we have opted to perform t-tests due to the simplistic nature of our study.

Elashoff (1981) makes a compelling argument that computing t-tests on multiple iterations for each point in time may lead to inconsistent and redundant results. Which is why we use t-tests only on aggregated values for a given experiment, i.e. only on 'Summary measures'.

To draw any meaningful inferences from the data and to analyse the validity of our

PVFR cooling control strategy over VFR and static cooling control strategies we need to employ a broad scope of testing. This broad scope should give an insight into the limitations of the cooling control strategies. This is why we have simulated the different cooling control strategies in alternate topological configurations, such as to test the upper/maximum limit of the control strategies in both the number of server racks and the server rack power densities for a fixed coolant pump capacity.

Efficacy of an intervention can be evaluated by analysing the data using t-tests (Elashoff, 1981), t-tests are employed as a method of data analysis to ensure only real differences between the intervention, subject and control are identified. Any randomness becomes apparent when analysing the data for p values using the predetermined statistical design. T-tests are based on studying variances in and between the sample group, rather than analysing only mean values (Rice, 1989).

6.3.2 Statistical design

Mason et al. (2003) states that for any statistical testing, there must be two specified hypotheses in order to test the effect of an intervention on the results. These hypotheses are the null and alternate hypotheses, denoted as H_0 and H_a respectively. The hypotheses both serve a different purpose, H_0 is used to indicate there will no experimental effect by defining parametric values. H_a on the other hand indicates the hypothesis to be tested, i.e. it specifies values in the experimental range.

In our study,

- H_0 is that, there will be no decrease in power consumption of coolant pumps when the PVFR control strategy is implemented over the VFR control strategy in LSDC's.

- H_a is that, there will be a statistically significant decrease in power consumption of coolant pumps when the PVFR control strategy is implemented over the VFR control strategy in LSDC's.

The terms used in any statistically designed experiment according to Mason et al. (2003), which are applicable to our study are as follows,

1. The Confidence Interval (CI), expressed as a percentage, is the probability that the results obtained will be applicable for given number of cases.

2. The Critical Value also known as 'p'value (p), is a cut-off value above which the alternate hypothesis is rejected. It is a function of the alternate hypothesis and confidence interval. Simply put, the smaller the p value, more pronounced is the alternate hypothesis. This implies a strong correlation between the intervention and the observed effect.

3. Degrees of freedom (n) is the number of independent values that can be assigned to determine a statistical distribution. Simplifying, the number of samples in a study - 1.

4. t-statistic value (t(n)), is a function of degrees of freedom and represents the measure of difference between the observed value and predicted value of an intervention.

In our study the p value is taken to be 0.5 with the CI value of 95%. This is the standard for all statistical studies employing t-tests (Mason et al., 2003). Finally, it is critical to define the type of t-test we are employing (Rice, 1989). In our study, as we cannot predict if there will be an increase or decrease in power consumption of the PVFR control strategy over VFR control strategy, it is suitable to employ a two-tailed t-test (Box et al., 1978). A two-tailed test represents the differences in either directions i.e. either increase or decrease in the measuring variables.

Non-reproducibility crisis for a study

A big issue with publishing scientific results in journals and conferences is the crisis in reproducibility of a study as ascertained by Lawrence and Lin (1989). They proposed two fundamental guidelines based on a quantitative metric designed to indicate if a study was easy to reproduce. The metric used a broad way of randomly sampling data to solve deterministic problems which usually arise in mathematical fields. The guidelines are as follows,

1. If the data points are profusely diffused the study is likely to be non-reproducible.
2. If duplicate values are used as distinct readings rather than being used as replicates the study is likely to be non-reproducible.

Solution

We have incorporated the guidelines proposed by Lawrence and Lin (1989) and the argument made by Treasure and MacRae (1998) into our study. This is why we have taken a few steps to make our results more robust, thereby avoiding the non-reproducibility crisis and to assure the inferences made on the results are accepted by the scientific cohort. A couple of points mentioned below help to illustrate the measures we have taken,

1. Raw data along with configuration and load files are unaltered and the URL for downloading the data is available in the Appendix A. This measure leads to greater transparency, which eliminates falsifying data points used in the study.
2. Using t-tests eliminates ambiguity in the data, which minimises any duplicate values having an effect on the results.
3. We have limited the number of measuring variables to only one, the power consumption. This measure safeguards any randomness that may have been

unknowingly introduced during experimentations has an imperceptible effect on the results.

6.3.3 Analysis

The data gathered from the simulations are neither dependent on each other, nor do we know the variance of the samples. Hence we need to use an unpaired two sample t-test assuming unequal variance for analysis (Rice, 1989). Reporting the results of the data analysis according to Jedlitschka and Pfahl (2005) is done by using Means, t-values and p-values. All of these terms have been explained in the previous statistical design section of this chapter.

We conducted an unpaired two sample t-test assuming unequal variance to compare the total power consumption of cooling pumps in PVFR and Static control strategies for data analysis. There was a statistically significant difference in the scores for PVFR (M=6209.55) and Static (M=17018.37) control strategies with conditions; t(11) = -49.93, p = 2.55 e^{-14}. These results suggest that there exists strong evidence to support that the PVFR control strategy has a considerably lower power consumption over the static control strategy. Specifically, our results suggest that the PVFR control strategy consumes approximately 63.51% less power to cool the server racks over a static control strategy when the mass flow rate is set to 6 kg/s.

An unpaired two sample t-test assuming unequal variance was conducted to compare the total power consumption of cooling pumps in PVFR and VFR control strategies. There was a statistically significant difference in the scores for PVFR (M=6209.55) and VFR (M=16530.79) cooling control strategies with conditions; t(12) = -50.54, p = 2.35 e^{-15}. These results suggest that there exists strong evidence to support that the PVFR control strategy has considerably lower power consumption over the VFR control strategy. Specifically, our results suggest the PVFR control strategy consumes approximately 62.44% less power to cool the server racks over the VFR control strategy.

An unpaired two sample t-test assuming unequal variance was conducted to compare the total power consumption of cooling pumps in VFR and Static control strategies. There was no statistically significant difference in the scores for VFR (M=16530.78714) and Static (M=17018.37429) control strategies with conditions; $t(12) = -2.09$, $p = 0.0576$. These results suggest that there exists no evidence to support that the VFR control strategy has any noticeable effect over the static control strategy when it comes to power consumption by cooling pumps. Specifically, our results suggest that the VFR control strategy consumes statistically the same power to cool the server racks over a static control strategy when the mass flow rate is set to 6 kg/s.

What do these results mean?

As evident from the results, if the PVFR control strategy is implemented over VFR or Static control strategies, data centre operators can expect to see at least a 62% reduction in power consumption by coolant pumps in liquid cooled data centres. The proves the efficacy of our proprietary PVFR control strategy over existing VFR and static flow rate control strategies.

Limitations of the study

T tests fail if the sample size (N) is not large enough, the rule of thumb for sample sizes while measuring a single variable is $N-1 > 5$, as stated by author Lachin (1981). The p value or probability value, which is used to determine if there exist a correlation or causality is taken to a statistically deterministic extent of $p < 0.5$. These guidelines when strictly followed ensure validity of a study. By collecting a large enough sample size, and measuring only a few variables, we can ascertain the correctness of inferences made. We have used data gathered across various points spread throughout the daisy chained server racks, to collect temperature readings to minimise the problems illustrated above.

However, as we have shown in this chapter the p value we have calculated far recedes the recommended 5% chance. The 5% chance measure ensures with any experimentation there exists a probability of the results being obtained through randomness. The chance measure is used to ascertain the efficacy of intervention itself rather than being by chance. Which is why there exists a general rule that p values « 0.05 show a strong evidence of the intervention working for all possible scenarios. i.e. p«0.05 states a strong argument can be made for the intervention being sound statistically (Rice, 1989).

6.4 Discussion

In this section, we investigate the results outlined from the analysis and discuss potential areas of PVFR implementations. We also calculate any cost savings that will occur if the PVFR cooling control strategy is implemented. This section also examines the correlations from a theoretical perspective to try and reason why exactly our PVFR cooling control strategy performs better against existing solutions i.e. VFR and static. We will use a logical reasoning methods known as 'Proof by contradiction' and 'proof by contrapositive' to test our hypothesis (Reeves & Brewer, 1980).

As CPS use physical control mechanisms, it is necessary to discuss the types of hystereses and the impact they have on cooling systems. The performance of the cooling systems also varies due to these hystereses.

6.4.1 Sensitivity to Hystereses

Hysteresis is the lag that occurs between a physical input to the desired output. There may be many such lags that occur within a thermal cooling system. These add up to become a performance issue. This is especially true in large complex systems, like our proposed system.

Some cooling control strategies are more prone to hysteresis than others, so for our discussion we have chosen the PVFR and VFR cooling control strategies. These two strategies were chosen because, there was no statistical difference between the static and VFR cooling control strategies. The most important factor for selecting the VFR cooling control strategy is the reason because it was modelled after information we could obtain for the dynamic flow AI managed control system Google which has already implemented (Gao, 2014).

The VFR cooling control strategy proves to perform similarly in terms of dynamic flow control systems while the static flow is a legacy implementation. The primary sources for impact on performance due to hysteresis are identified as follows.

1. Lag between fluctuating flow rate by actuating coolant pumps and desired flow occurring i.e. lag between controller and physical components.

2. Lag between controller polling temperature sensors of all the server racks and actuating coolant pumps i.e. lag between physical components and controller.

6.4.2 Sensitivity to response time between controller and physical components

This is the time it takes for physical components in the thermal network receiving input to the time it takes to performing the expected action. As an example, we will discuss the hysteresis between actuation of the pump to the desired effect of achieving intended flow rate.

The time between these components vary significantly with the design choice made for actuators, as examples hydraulics, solenoids etc. we shall take 'T_a' as the time required to perform the actuations, and 'T_{cp}' time required for the coolant pump to provide the expected flow rate. This is because pumps cannot provide coolant flow instantaneously.

Therefore, the equation to represent the total hysteresis is as follows,

$$Total\,Hysteresis\;\; T_h = n * (T_a + T_{cp}) \tag{6.1}$$

Where, T_a = time required to perform the actuations, T_{cp} = time required for the coolant pump to provide the expected flow rate and n = number of times it is performed.

Let us apply the technique of proving by contrapositive to prove our hypothesis that, the PVFR cooling control strategy has a lower latency of hysteresis when compared to VFR cooling control strategy. Let us assume the contrapositive of the hypothesis, which is that the VFR cooling control strategy has a lower latency when compared to the PVFR cooling control strategy. The principle involved in achieving a dynamic flow using a VFR cooling control strategy is that it actuates pumps based on temperature ranges.

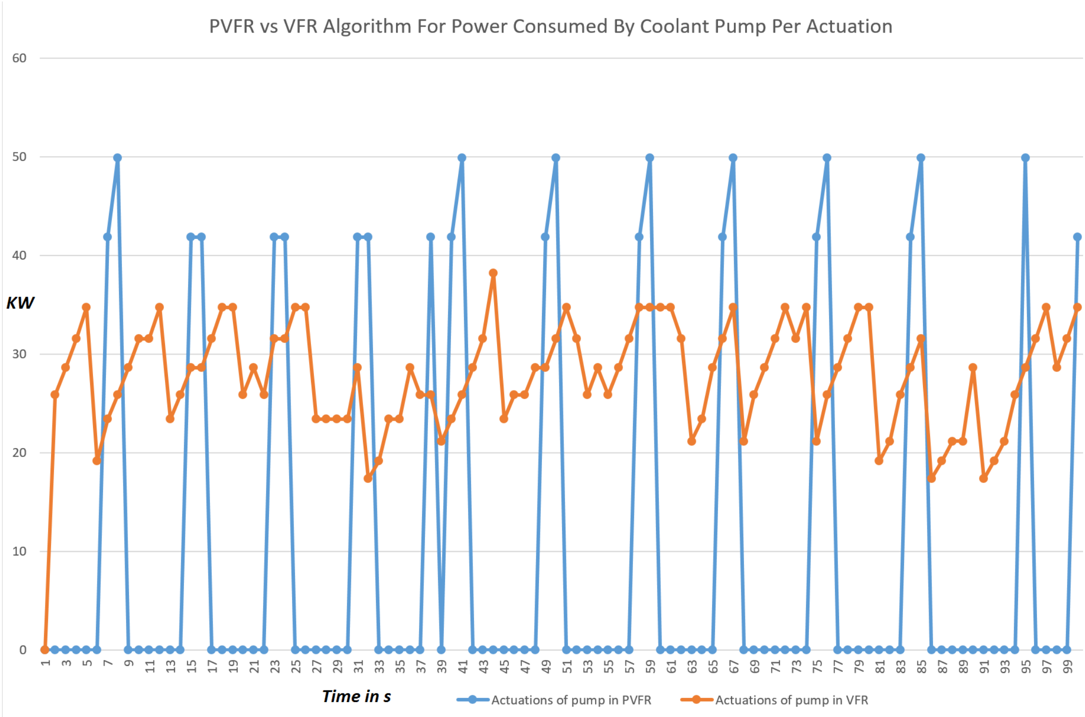

Figure 6.12: Power consumed by a single pump to illustrate the no. of actuations in the working of PVFR and VFR cooling control strategies

Figure 6.12 shows the number of actuations each cooling control strategy performs for a single pump. The VFR cooling control strategy performs approximately ten actuations per ten second while PVFR performs approximately only five. This is an actuation ratio of 2:1, from this we can deduce PVFR should perform twice as better. But this is not a good inference as this only takes into account the T_a time of actuation. T_{cp} is a variable and increases with increase in flow rate, this is because the molecules accelerate based on the Newton's law of inertia with time. So let us assume for simplicity it takes equal time 'x' to increase the mass flow rate for every unit, in the figure 50 kW represents a mass flow rate of 10 kg/s, therefore approx. 10 kW of power consumed equates to a mass flow rate increase of 2 kg/s. From the Figure 6.12 we can also note that it takes approx. 0.2 s to increase the flow rate by 1, computed using power consumed by pump.

T_{cp} (VFR) = 10 * 0.08 = 0.8s for every 10 seconds of operation and T_{cp} (PVFR) = 5*0.2 = 1s for every 10s of operation. which is a ratio of 1:1.25, solving equation 6.1 using the two ratios we get,

T_h (VFR) = 2x+1x = 3x and T_h (PVFR) = 1x+1.25x = 2.25x   Where,   x= some unit of time.

We cannot solve for x, since we do not enough information to solve it completely. But we do however have an answer for the total hysteresis incurred by the two cooling control strategies, The PVFR has a 25% hysteresis improvement over VFR. This is contrary to our assumed hypothesis thereby proving our original hypothesis that PVFR has lower hysteresis i.e. lag when compared to VFR under identical conditions.

6.4.3 Sensitivity to response time between physical components and controller

Fozdar, Parkar and Imberger (1985) performed an experimental study to determine the response rate of a thermal sensor (SBE-3). The manufacturer specified response time was 70 ms, the researchers found that a time constant of 67 ms was required for a step response. This is in close agreement with the manufacturers specified times and as such they are used in our calculations.

This is a simple calculation as it is a sum of all the sensors being measured i.e. the number of server racks. This is a linear relation since the more number of server racks equates to more of a lag. In our test configuration of 10 server racks, the lag equated to 10*70 ms = 0.7 s. Which is we are recommending a *polling rate for the controller to be set to half the response rate*. The polling rate in our example was set to 0.5 s because the response rate was approximately 1 s.

This lag is manageable and is the same for all cooling control strategy implementations. It is a function of the number of server racks being measured. This lag can be minimized by approximating the temperatures of servers by measuring only a few server racks rather than all of them. This is not recommended as it may lead to unexpected consequences and erroneous timings in coolant pump actuations. This is because the temperature of the server racks is a function of the load placed on them and as a result they can vary with dynamic load.

6.4.4 Sensitivity to Thermal Loads

Conventional wisdom dictates larger thermal heat loads require increased cooling i.e. they are linear (Chandler, 1987). However, in practice, especially when considering server racks where high thermal output is expected in a confined space, this assumption leads to erroneous results.

Let us use proof by contradiction to prove our hypothesis, that higher thermal loads have a non-linear impact on cooling demand. We shall start by assuming the inverse of our hypothesis or conventional wisdom as true i.e. there is a linear relationship between thermal heat load and cooling demand. Referring to the results seen in Figure 6.11, we can determine a thermal load of 600 kW/s consumes 35.21 kW by the pumps. A thermal load of 300 kW/s consumes 17.01 kW by the pumps, from Figure 6.6, and a thermal load of 1000 kW/s consumes 41.34 kW by the pumps from Figure 6.9. These results are for a static flow rate cooling control strategy which ran for 10 minutes.

Mapping these values on a spreadsheet and fitting a curve using linear regression we get the following curve illustrated in Figure 6.13. Since a linear curve is not obtained we can deduce there exists a non-linear relationship between the thermal load placed on the system and cooling demand. This statement is contrary to our assumed hypothesis of conventional wisdom, thereby proving our original hypothesis that higher thermal loads have a non-linear impact on cooling demand. This was also evidenced by researchers Arguello-Serrano and Velez-Reyes (1999) while designing a control system for air cooling data centres.

This unconventional cooling demand leads us to explore an intriguing question, which is to uncover from a physics stand point the reason behind such a strange outcome. When considering the mechanics of fluids, fundamentally heat exchange is a function of the difference between the temperature of the cooling fluid and the heat source (Shames & Shames, 1982).

Since our configuration requires servers to be daisy chained, the differences in temperatures become more pronounced with the addition of more number of server racks. Considering a ratio of cooling power consumed to the server rack power density, we see an anomaly with the increase in power consumption by pumps to maintain the same server temperatures even at a lower thermal load of 600 kW/s.

Therefore, to minimise the negative effects of this non-linear sensitivity to thermal

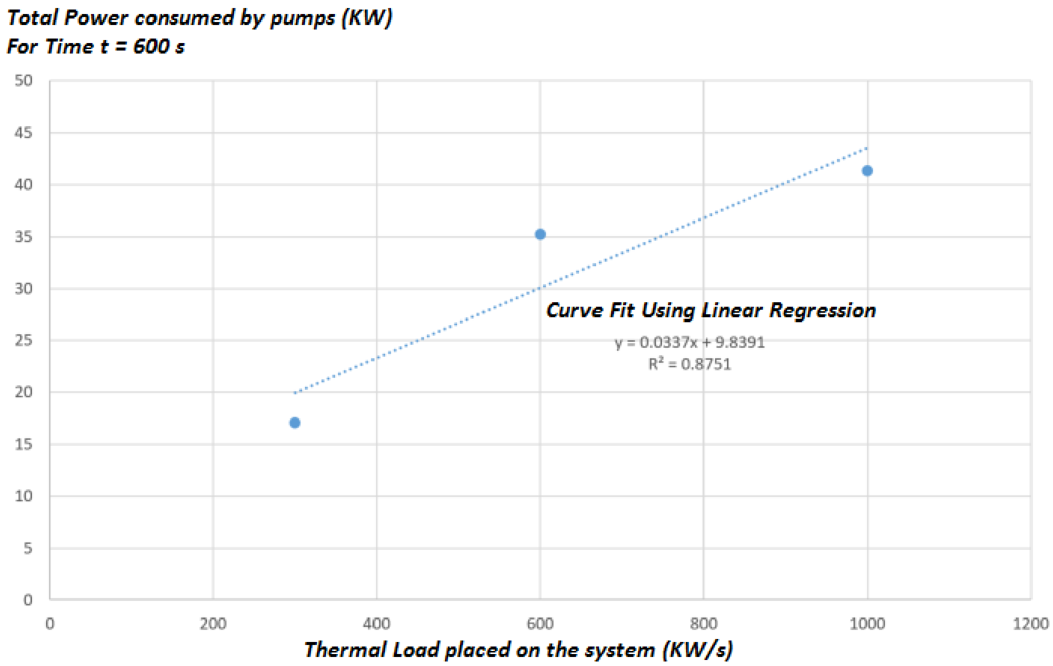

Figure 6.13: Disproving Assumptions of Linear regression on heat load

.

heat load, we recommend embracing modularisation of data centres. This means limiting a row of server racks to a *ratio of 1:1 in terms of number of servers to the maximum mass flow rate of cooling pump, assuming the heat capacity of the coolant is equal to that of water at 4.18 kJ/kg*. For example, ten servers for a maximum flow rate of 10 kg/s. This implies a greater number of server racks can be used if the properties of coolant can be enhanced. We also recommend data centre operators use a *ratio of 10:1 in terms of the power density of a 42U server rack to the mass flow rate of cooling pump, assuming the heat capacity of the coolant is equal to that of water at 4.18 kJ/kg*. For example, a 100 kW/s per rack power density for a 10 kg/s maximum mass flow rate coolant pump.

6.4.5 Impact of the PVFR cooling control strategy

Financial savings

Cooling in data centres account for approximately 40% of total power consumed (Capozzoli & Primiceri, 2015). So in the best case, we can expect the PUE to increase by 30.8%; this imputes a cumulative direct cost saving of *$1.386 billion* for the 61 TWH of electricity consumed by data centres in the US alone (R. Brown et al., 2008). This figure does not take into account infrastructure and other indirect cost savings such as modularisation, server utilisation etc. Therefore, the total cost savings are much higher.

This could make the PVFR system implementation feasible and financially attractive for retrofitting existing data centres. The costs for installation in such an event is minimal, since the system does not use any complex machine learning algorithms nor does it require large amount of computing resources. The post-installation benefits could return an investment in the first year or even the first quarter, depending on certain criteria like size, density, thermal output etc.

Putting the results into context

Although the energy savings equate to a PUE saving of 30.8%, this should be put into context to determine the true cost savings. Traditional liquid cooling systems use a chiller or cooler to cool the temperature of circulating coolant, these systems are the single largest energy consumers in any cooling system. Which is why Google have moved to a free air cooler design, where by the coolant is cooled through natural convention of air.

This design allows to eliminate chillers or coolers and increase efficiency through utilising natural processes for heat dissipation (Google, 2016a). Google is pioneering innovative cooling technologies, with a recently opened data centre using ocean water for cooling the site in its entirety (Google, 2017b).

Putting this into perspective for our technology we can deduce that our system needs a water-water heat exchanger. At least, in this instance we need such a system to rapidly cool the coolant to achieve the temperature differences needed for operation. This should change with further research and calibration to allow for free air coolers.

Therefore, the system would use approximately twice as much energy. This is because we have synchronised the water cooling system to consume as much energy as the cooling pumps. In other words, *Real PUE Reduction = 30.8/2 = 15.4%*. This calculation assumes the worst case, and without any proper calibration. Google's Deepmind based machine learning implementation of automated control systems have published total reductions in cooling power consumed by 40% and reduction in PUE by 15% (Google, 2017a). While this is a significant decrease in power consumption with Google boasting the world's best fleet wide PUE at 1.1 (Google, 2016a), there is still room for improvement.

Reduction in green house gas emissions and cost savings for carbon mediation

Sims et al. (2003), evaluated the carbon footprint of current power generation methods. They estimated 151 g of greenhouse gases are released into the atmosphere per KWH of electricity, averaged from all sources such as renewables, coal, oil etc. Furthermore, they calculated the costs of carbon mitigation, it was anywhere from 9-25 cents per KWH depending on the type of power plant used.

The power savings mentioned above using the calculations proposed by Sims et al. equate to approx. 18.788 TWH of electricity, which translates to 18,788,000,000 KWH. This power reduction saves 2,836,988 tons of greenhouse gas emissions or *just under 2.84 million tons of green house gases per year! being released into the air we breathe*. The cost to mediate this is around $3,193,960,000 @ 17 cents/KWH. or approx. $3.2 billion per year. While these figures are calculated for an adoption rate of 100%, it is impractical and a 10% adoption rate would yield savings of approximately $319 million and 284,000 tons of greenhouse gas emissions per year. These savings are significant while considering even a low adoption rate of 10%.

6.4.6 Strengths and limitations of the PVFR cooling control strategy

The advantage of using a PVFR approach is that it forces data centre operators to modularise or zone their infrastructures which further increasing efficiency. This is because zones can be turned off during periods of low demands and offers better load scheduling to keep zones at a higher server utilisation rate.

It also allows for increased flexibility in terms of speed in scaling, as this approach incorporates decoupling between the different components of modules. This allows data centre operators to add or subtract components of modules or the modules themselves based on demand forecasts.

The disadvantage of our cooling control strategy over a model similar to that of Google's is that, our approach is optimised for only one type of application, i.e. high thermal output applications. Google's approach uses generalised intelligence, which they have leveraged to offer better services in other fields. Our improvement over Google's should not be taken as absolute, rather it should be considered as an alternate approach to optimise data centres.

6.5 Summary

Section 6.1 details the experimental setup used to conduct the experiments for executing and gathering data from the DynaCool simulation software. Section 6.1 also expresses the different experimental setup used to gather a comprehensive set of simulation data. Section 6.2 illustrates the data gathered visually through graphs for the different experimental setups. Section 6.3 discusses the data analysis method used and the statistical design used in our study. Section 6.4 exposes the sensitivities of the PVFR cooling control strategy to various factors and the impact it has on the data centre industry. The impact shows the financial and environmental benefits of implementing the PVFR cooling control strategy to reduce power consumption of coolant pumps. Section 6.4 discusses the strengths and limitations of the PVFR cooling control strategy.

Chapter 7

Conclusion

This chapter summarises the results obtained by simulating efficient liquid cooling for current and next generation Large Scale Data Centres (LSDC's). The contributions made in this thesis overarch a broader research domain of efficient cooling of thermal dissipations in high heat load applications.

Section 7.1 provides a summary to this thesis as a whole and describes the observations and insights obtained in previous chapters of this thesis. Section 7.2 focuses on summarising the answers to the research questions identified previously in the literature review chapter of this thesis (refer Section 2.1.2). Section 7.3 details the contributions made while undertaking the research, the optimal cooling control strategy and its corresponding recommended data centre topological configuration. Section 7.3 also discusses some of the directions for future work that could be undertaken and the potential avenues for further research. Finally, Section 7.4 outlines our final words regarding this research and the thesis.

7.1 Summary

We investigated several different methods of cooling, viz. air, liquid and hybrid, all of which have been implemented by high technology giants like Facebook (Facebook, 2010) and Google (Google, 2016a). Although all the systems have calibre, liquid cooling provided better opportunities for research contribution as liquid cooling systems in LSDC's are still an emerging technology (Dai et al., 2014). For more information on why liquid cooling systems were chosen and the limitations of other approaches please refer the literature review chapter (refer Section 2.4.2). We inspected the two different principles involved in liquid cooling systems, which are static flow rate and dynamic flow rate. The merits of both approaches were studied, which led to the conclusion that dynamic flow rate liquid cooling systems are superior (refer Section 2.5).

We scrutinised Google's approach of liquid cooling and modelled a logic controller to develop a cooling control strategy known as Variable flow rate(VFR) (refer Section 5.2), for dynamic cooling of server racks based on Google's AI liquid cooling control strategy (Gao, 2014). Although their AI implementation provided a substantial decrease in Power Usage Effectiveness (PUE) (Google, 2017a), there was a key limitation that it needed large amounts of data and computing resources for it to function effectively, as noted by Gao (2014). This is because it uses a generalised machine learning algorithm, and is not specialised for any particular application, which led us to advance the field by developing a novel implementation to optimise the dynamic control strategy.

We hypothesised that using pulses of coolant rather than a continuously varying flow rate would yield better efficiencies. The hypothesis aimed to exploit the increases in thermal tolerances of micro-processors and the enhanced efficiency of coolant pumps when running at full load to reduce power consumption. Choi et al. (1994) studied and patented this energy saving approach in small scale applications such as handheld battery powered light emitting diodes by proving flickering the lights extends battery life. However, for high power applications the theory was untested and proved to be a novel approach in cooling LSDC's.

We opted to use simulation software to test the hypothesis instead of a physical LSDC due to the constraints explained in the literature review chapter (refer Section 2.6.4), but obtaining such a simulation software which met our requirements proved to be challenge. We viewed the lack of such a software an opportunity rather than a setback, which is why we developed a proprietary simulation software called 'DynaCool' and made it Open Source. Open Source Software's help other researchers in the field and promotes further development in the community (refer Appendix A.1 for the URL/link to download the software).

The DynaCool software was validated in order for the data obtained from the software to be accepted by the scientific community. We validated the software against figures published by data centre operators and using CFD software to estimate the interaction of molecules. By using a sound mathematical approach, such as using CFD calculations we were able to validate the DynaCool software accurately within the accepted error margin of 5%. This means the DynaCool system accurately simulates the functions of a physical data centre, with the ability to test different liquid cooling control strategies.

The cooling control strategy we developed from our initial hypothesis is known as Pulsed Variable Flow Rate (PVFR), whose merits and weaknesses have been outlined previously in the experimental results chapter (refer Section 6.4.6) and discussed later in this chapter (refer Section 7.2.2). The PVFR cooling control strategy can be implemented by retrofitting existing data centres at a cost that would not inhibit adoption. In our testing, we found that it would provide a 'Real PUE reduction' of approximately 15.4% or higher. A detailed summary of the contributions made while undertaking this research can be found in later in this chapter (refer Section 7.3).

7.2 Answering the research questions

7.2.1 Research question 1

The first research question is, What are the key functional and non-functional requirements for cooling systems in next generation LSDC's?

The cooling requirements of next generation LSDC's were outlined as part of our systematic literature review (refer Section 2.6.2). To summarise, next generation LSDC's will place a cooling demand of 100 kW per server rack on cooling systems, while the physical footprint of the data centres remain constant thereby promoting increasing density between individual server racks. The final requirement is to reduce the carbon footprint of cooling systems by either decreasing power consumption or using more efficient dynamic cooling approaches.

We identified the functional and non-functional requirements in the system design chapter of this thesis (refer Section 4.1). Requirement analysis was done using the principles involved in the Requirements Engineering (refer Section 4.1.3). Finally, these requirements were used to design the DynaCool system using Model Driven Engineering (refer Section 4.2).

7.2.2 Research question 2

The second research question is, How can we design more efficient cooling systems for next-generation LSDC's such that key functional and non-functional requirements are met?

The answer this question led us to explore and evaluate the different cooling approaches and develop the novel Pulsed Variable Flow Rate (PVFR) cooling control strategy. The summaries of this approach is illustrated in the following sections.

Cost and greenhouse gas emissions savings

The reduction in power consumption allows the data centre industry to save approx. $1.386 billion per year if the PVFR cooling control strategy is implemented in all data centres. Such an event is unlikely for current data centres but is highly likely to be incorporated into next-generation data centres and current generation data centres from high technology giants like Google. Such a low adoption rate is one of the reasons why the PVFR cooling control strategy is designed for future LSDC's cooling demands. Furthermore, the environmental impact of the PVFR cooling control strategy is that it saves approx. **2.84 million tons of greenhouse gases per year being released into the air we breathe**. The cost to mediate 2.84 million tons of carbon emissions is $3,193,960,000 @ 17 cents/KWH. or approx. $3.2 billion per year based on the figures published by Sims et al. (2003), which is an indirect cost saving for data centre operators. Considering an adoption rate of 10%, which is highly likely, we can expect savings of $319 million and 284,000 tons of greenhouse gas emissions per year.

Strengths and limitations of PVFR

Summarising the merit of the approach, it was found to be more efficient than Google's implementation utilising fewer computing resources (refer Section 6.3). Furthermore, from the data analysis (refer Section 6.3.3) we can conclusively establish a strong relation between the efficiencies of three different control strategy implementations. These represent the models of variable flow control in existing LSDC control systems and our proprietary dynamic cooling PVFR approach. The efficacy of the PVFR control strategy using our novel PVFR algorithm is at least 25% better in the worst case and 77% better in the best case, over the other two cooling control strategy implementations.

The PVFR control strategy performed best when the power density of server racks was 30 kW per server rack, which is the density of current LSDC's (R. Schmidt et al., 2005) and the number of daisy chained server racks was 10. Increasing either the number of server racks or their densities led to diminishing performance gains, but still outperformed the other control strategies by a significant margin. The upper limits of thermal load for cooling pumps with maximum 10 kg/s was found to be at 1 MW, when the number of servers were 10 thereby imputing a power density of 100 kW per server rack, which are the predicted densities of future LSDC's (Musilli & Ellison, 2012).

The feedback loops of the VFR implementation for increasing efficiencies failed to operate at an effective rate at these higher densities making it unsuitable for future LSDC cooling control systems. Static and PVFR performed at an acceptable level at these power densities. Further calibration attempts may prove useful in increasing efficiencies as these systems have not been calibrated properly and have been simulated for first time use only.

The bottleneck of the PVFR implementation is in the number of server racks daisy chained, rather than increasing power densities, this is due to the laws of thermodynamics which was alluded to later in the previous literature review chapter of this thesis (refer Section 6.4.6). This is a key limitation and a potential avenue for future research, and as such it is discussed later in this chapter (refer Section 7.3.4).

7.3 Contributions

By undertaking the research to enhance the efficiency of liquid cooling in LSDC we make the following contributions.

7.3.1 DynaCool

The DynaCool simulation software is an application to test different control strategies for liquid cooling in data centres and is based on similar CFD software, such as CloudSim (Buyya et al., 2009). The key differentiating factor for DynaCool is that it is designed to provide data on liquid cooling control strategies for stakeholders (refer Section 4.1.1, unlike CloudSim which is limited to air cooling.

The developed DynaCool simulation software dispenses accurate and valid data for analysts to designate optimal control logic in their data centre thermal management systems. DynaCool, is the first CFD simulation software offering such flexibility for stakeholders to optimise and test their cooling control strategies. This is because current commercial systems like CloudSim can only simulate air cooling, which is a serious limitation for data centre operators to plan future LSDC designs and implementation strategies.

The decoupling of control logic from the data centre models is a compelling argument to support the merits of the DynaCool system. DynaCool can also use dynamic data centre models to represent different requirements for different infrastructures. This feature allows analysts to test a specific control strategy for their whole fleet of data centres. The data allows operators to implement custom tailored solutions to achieve the highest possible energy reduction.

7.3.2 A more efficient cooling control strategy

The proposed Pulsed Variable Flow Rate (PVFR) control strategy for cooling in high heat load applications like LSDC's is optimally suited to reduce energy consumption by cooling pumps. This reduction is significant and has compelling evidence for implementation in current and future LSDC's. The control strategy is also distributed enabling zoning or modularisation of data centres. This will have a direct impact on server usage, the reasoning being a higher usage leads to lower costs in cooling.

IT load or demand is a real time varying component, according to Gao (2014), Google's approach is to use complex machine learning algorithms to predict future usage trends based on several factors like peak usage curves. This is a highly complex and resource intensive endeavour. Our approach is much simpler, and requires in comparison trivial resources. We have exhaustively tested this approach against Google's approach modelled using the fundamental principle of continuously varying the flow rate. This model in our opinion is an accurate representation of the working of any such approach in principle.

Performance metrics like response times were measured taking into account different forms of hystereses that will be induced during operation. In all measures the PVFR approach proved to perform better and produce the *highest reduction or a real PUE of 15.4% or higher*. This is slightly better when compared to Google's reduction of 15% (Google, 2017a). Although not by much, this is based on direct reductions, not taking into account indirect costs like computing resources, installation, maintenance costs etc.

7.3.3 Optimal Data Centre Topology Model

We have provided recommendations for current and future data centre operators to provide an optimal layout of data centre infrastructural components for achieving the lowest PUE (refer Section 6.4). These recommendations are as follows,

- The polling rate of the sensors from the thermal system controller should be set to half the response rate of the pumps. or in other words, poll the sensors twice as fast as the response rate of the pumps to ensure smooth operation.
- Adopt a ratio of 1:1 in terms of the number of servers to the maximum mass flow rate of the cooling pumps, assuming the heat capacity of the coolant is equal to that of water at 4.18 kJ/kg
- Adopt a ratio of 10:1 in terms of the power density of a 42U server rack to the mass flow rate of cooling pump, assuming the heat capacity of the coolant is equal to that of water at 4.18 kJ/kg

The recommendations are based on the data analysis methods recommended by Chandler (1987). The data was gathered by simulating different configurations of data centre layouts in the DynaCool software (refer Section 6.2).

The recommendations we have proposed provide a holistic overview for analysts to design and implement an optimal control strategy. It also useful to note that both the simulation and subsequent analyses were done assuming a stochastic IT load. While this is not representative of real world demand, it is indicative of an exhaustively aggregated data set. Therefore, an analysis done using such data sets will prove to be valid for all implementations within an acceptable range of error.

7.3.4 Future Work

There are multiple possible research and development avenues for future contributions in this field. Some of the opportunities that can be pursued by future researchers are as follows:

Further development of DynaCool

The current version (V2.1) of the DynaCool software has no Graphical User Interface (GUI) and having such a feature could prove to be a useful contribution. Another possible area for contribution is to add secondary and auxiliary CFD equations for simulations.

The current version only looks at the primary interactions for CFD analysis. This is because fluid dynamical interactions are notoriously hard to predict. This problem is further exacerbated by the addition of heat exchangers. Time constraints for this thesis have led us to only look at the primary interactions as they are the main contributors for physical changes. Secondary and auxiliary interactions could prove to influence changes in the system. This is however just a theory and the interactions may not be noticed in practice and as such is an intriguing research problem.

Humidity calculations

Maintaining safe and acceptable humidity parameters is a critical aspect of ASHRAE's guidelines (ASHRAE, 2011) and our research has focused mainly to reduce power consumption of coolant pumps using efficient cooling control strategies. We have assumed humidity is not a concern since we are opting for a liquid cooling approach and not a hybrid one, where humidity is an essential factor for consideration. Further evaluation on our assumptions should be an interesting research avenue for future researchers.

Model Validation

The DynaCool system and the deduced models have been validated against real world data, but this is only from a few published sources, which represents only a small subset of the total data that can be obtained by real world implementation. Validation against a larger data set is the most important research direction for future researchers to pursue. The distributed nature of the architectures involved lead us to expect that the stated benefits can not apply in every single instance of implementation but proves to be better than the stated benefits in some cases and worse in other cases.

As stated earlier, there may be unaccounted interactions that may prove to influence different data centre configurations. Because of which the results may vary, and is why it would prove advantageous to validate the model against a varied range of different data centre configurations.

Data centre infrastructural components

The current DynaCool version only looks at the most prominent infrastructural components and fails to account for all the possible intricacies of these components. The lack of implementing all the possible infrastructure components is because in principle all the components perform similarly from the stand point of thermodynamics. for example, strictly speaking for thermodynamic processes a 3 kW heat load from a network switch is the same as 3 kW heat load from network drives. But in theory, these devices could perform differently as their surface areas may vary, which is a key component for heat exchange. In practice the thermodynamic behaviours may be small but could prove to affect other processes. Investigation and validation of the system accounting for these factors pose an interesting research problem.

7.4 Final words

Undertaking this research posed a daunting challenge at first, but through hard work and perseverance the challenges faced proved to be opportunities worth investigating. The goal we set out to accomplish was to reduce the PUE of LSDC's by employing efficient cooling control strategies and data centre topology designs. Our contributions of providing the first open source liquid cooling simulation software, *DynaCool* along with optimised data centre topologies and a novel cooling control strategy, Pulsed Variable Flow Rate (PVFR) can potentially help reduce greenhouse gas emissions from LSDC's.

The research done in this thesis should provide interested researchers all the tools required for further development and enhancements of PUE and testing the strategies on physical LSDC's. We have shown that the PVFR cooling control strategy reduces power consumption of coolant pumps significantly with a reduction in PUE of 15.4%. This reduction equates to saving approx. 2.84 million tons of greenhouse gas emissions per year and power savings of 18.788 TWH (Based on the data centre power consumption data of 61TWH in 2006 published by USA's environmental protection agency (R. Brown et al., 2008)). Realistically, considering an adoption rate of 10%, which is highly likely, we can expect savings of $319 million and 284,000 tons of greenhouse gas emissions per year.

As we stated in the beginning (refer Section 1.6), our three criteria for success have not only been met, but exceeded. We are ecstatic about the contributions made in such a short time frame for a 90 point thesis.

References

3M. (2016). Two-phase immersion cooling. Retrieved 11 December 2016, from http://www.3m.com/3M/en_US/novec/products/engineered-fluids/immersion-cooling/

Adams, K. (2015). *Non-functional requirements in systems analysis and design* (Vol. 28). Springer.

Alford, M. W. (1977). A requirements engineering methodology for real-time processing requirements. *IEEE Transactions on Software Engineering*, *3*(1), 60.

Allen, A. T. (1982). *Sas user's guide: basics* (Vol. 1). SAS Institute.

Allen, D. J. & Lasecki, M. P. (2001). *Thermal management evolution and controlled coolant flow* (Tech. Rep.). SAE technical paper.

Almoli, A., Thompson, A., Kapur, N., Summers, J., Thompson, H. & Hannah, G. (2012). Computational fluid dynamic investigation of liquid rack cooling in data centres. *Applied energy*, *89*(1), 150–155.

Aquatherm. (2015). Data center liquid-cooling systems with polypropylene pipe. *White Paper by Aquatherm*.

Arguello-Serrano, B. & Velez-Reyes, M. (1999). Nonlinear control of a heating, ventilating, and air conditioning system with thermal load estimation. *IEEE Transactions on Control Systems Technology*, *7*(1), 56–63.

ASHRAE, T. (2011). Thermal guidelines for data processing environments-expanded data center classes and usage guidance. *Whitepaper by ASHRAE TC*, *9*.

Askwig, J. Q., Butler, N. R., Huglen, K. M. & Sinda, S. M. (n.d.). Server cooling comparison direct liquid immersion vs. tradition air cooling.

Autodesk. (2013). Air, fluid flow, and thermal simulation of data centers with autodesk revit 2013 and autodesk bim 360. *White Paper by Autodesk*.

Avison, D. E., Lau, F., Myers, M. D. & Nielsen, P. A. (1999). Action research. *Communications of the ACM*, *42*(1), 94–97.

Bachmann, F. & Bass, L. (2001). Introduction to the attribute driven design method. In *Proceedings of the 23rd international conference on software engineering* (pp. 745–746).

Baliga, J., Ayre, R. W., Hinton, K. & Tucker, R. S. (2011). Green cloud computing: Balancing energy in processing, storage, and transport. *Proceedings of the IEEE*, *99*(1), 149–167.

Barroso, L. A., Clidaras, J. & Hölzle, U. (2013). The datacenter as a computer: An introduction to the design of warehouse-scale machines. *Synthesis lectures on* *computer architecture*, *8*(3), 1–154.

Basili, V. R. & Reiter, R. W. (1981). A controlled experiment quantitatively comparing software development approaches. *IEEE Transactions on Software Engineering*(3), 299–320.

Basili, V. R., Selby, R. W. & Hutchens, D. H. (1986). Experimentation in software engineering. *IEEE Transactions on software engineering*(7), 733–743.

Basili, V. R., Selby, R. W. & Phillips, T. (1983). Metric analysis and data validation across fortran projects. *IEEE Transactions on Software Engineering*(6), 652–663.

Basili, V. R. & Selby Jr, R. W. (1986). Four applications of a software data collection and analysis methodology. In *Software system design methods* (pp. 3–33). Springer.

Belady, C., Rawson, A., Pfleuger, J. & Cader, T. (2008). *Green grid data center power efficiency metrics: Pue and dcie* (Tech. Rep.). Technical report, Green Grid.

Bell, T. E. & Thayer, T. A. (1976). Software requirements: Are they really a problem? In *Proceedings of the 2nd international conference on software engineering* (pp. 61–68).

Bergman, T. L. & Incropera, F. P. (2011). *Introduction to heat transfer*. John Wiley & Sons.

Board, I. (1998). Ieee guide for developing system requirements specifications. In *Institute of electrical and electronics engineers* (Vol. 1).

Booch, G., Rumbaugh, J., Jacobson, I. et al. (1999). *The unified modeling language user guide.* Addison-Wesley Reading, MA.

Box, G. E., Hunter, W. G., Hunter, J. S. et al. (1978). Statistics for experimenters.

Brereton, P., Kitchenham, B. A., Budgen, D., Turner, M. & Khalil, M. (2007). Lessons from applying the systematic literature review process within the software engineering domain. *Journal of systems and software*, *80*(4), 571–583.

Brown, K. (1995). Remembrance of things past: layered architectures for smalltalk applications. *The Smalltalk Report*, *4*(9), 4–7.

Brown, R. et al. (2008). Report to congress on server and data center energy efficiency: Public law 109-431. *Lawrence Berkeley National Laboratory*.

Buyya, R., Ranjan, R. & Calheiros, R. N. (2009). Modeling and simulation of scalable cloud computing environments and the cloudsim toolkit: Challenges and opportunities. In *High performance computing & simulation, 2009. hpcs'09. international conference on* (pp. 1–11).

Capozzoli, A. & Primiceri, G. (2015). Cooling systems in data centers: State of art and emerging technologies. *Energy Procedia*, *83*, 484–493.

Chandler, D. (1987). Introduction to modern statistical mechanics. *Introduction to Modern Statistical Mechanics, by David Chandler, pp. 288. Foreword by David Chandler. Oxford University Press, Sep 1987. ISBN-10: 0195042778. ISBN-13: 9780195042771*, 288.

Chang, Y.-W. & Hsu, P.-Y. (2016). Investigating switching intention to cloud enterprise information systems: an analysis at the organizational level. *International Journal of Information Management*.

Chen, G., He, W., Liu, J., Nath, S., Rigas, L., Xiao, L. & Zhao, F. (2008). Energy-aware server provisioning and load dispatching for connection-intensive internet services. In *Nsdi* (Vol. 8, pp. 337–350).

Chen, T. Y. & Poon, P.-L. (2004). Experience with teaching black-box testing in a computer science/software engineering curriculum. *IEEE Transactions on Education*, *47*(1), 42–50.

Cho, H., Jung, D., Filipi, Z. S., Assanis, D. N., Vanderslice, J. & Bryzik, W. (2007). Application of controllable electric coolant pump for fuel economy and cooling performance improvement. *Journal of engineering for gas turbines and power*, *129*(1), 239–244.

Choi, R. S., Kong, C. K. & Choi, O. (1994, May 17). *Battery-powered flashing superluminescent light emitting diode safety warning light.* Google Patents. (US Patent 5,313,187)

Chu, R. C. (2003). The challenges of electronic cooling: past, current and future. In *Asme 2003 international mechanical engineering congress and exposition* (pp. 619–630).

Chung, L., Nixon, B. A., Yu, E. & Mylopoulos, J. (2012). *Non-functional requirements in software engineering* (Vol. 5). Springer Science & Business Media.

Clark, J. (2013). Facebook's first data center drenched by actual cloud. Retrieved 07 December 2016, from `http://www.theregister.co.uk/2013/06/08/facebook_cloud_versus_cloud/`

Cochran, W. G. (2007). *Sampling techniques*. John Wiley & Sons.

Committee, I. C. S. S. E. S. & Board, I.-S. S. (1998). Ieee recommended practice for software requirements specifications..

Conservative & coalition government of the UK, L. D. (2016). 2010 to 2015 government policy: energy demand reduction in industry, business and the public sector. Retrieved 07 December 2016, from `https://www.gov.uk/government/publications/2010-to-2015-government-policy-energy-demand-reduction-in-industry-business-and-the-public-sector/2010-to-2015-government-policy-energy-demand-reduction-in-industry-business-and-the-public-sector`

Cronin, P., Ryan, F. & Coughlan, M. (2008). Undertaking a literature review: a step-by-step approach. *British journal of nursing*, *17*(1), 38.

Dai, J., Ohadi, M. M., Das, D. & Pecht, M. G. (2014). *Optimum cooling of data centers*. Springer.

David, M. P. & Schmidt, R. R. (2014). Impact of ashrae environmental classes on data centers. In *Thermal and thermomechanical phenomena in electronic systems (itherm), 2014 ieee intersociety conference on* (pp. 1092–1099).

Deacon, J. (2009). Model-view-controller (mvc) architecture. *Online][Citado em: 10 de março de 2006.] http://www. jdl. co. uk/briefings/MVC. pdf*.

De Groot, S. R. & Mazur, P. (2013). *Non-equilibrium thermodynamics*. Courier Corporation.

Dubinin, V. & Vyatkin, V. (2007). On definition of a formal model for iec 61499 function blocks. *EURASIP Journal on Embedded Systems*, *2008*(1), 1–10.

Dumas, M. & Ter Hofstede, A. H. (2001). Uml activity diagrams as a workflow specification language. In *International conference on the unified modeling* *language* (pp. 76–90).

Dwyer, V., Franklin, A. & Campbell, D. (1990). Thermal failure in semiconductor devices. *Solid-State Electronics*, *33*(5), 553–560.

Dybå, T., Kampenes, V. B. & Sjøberg, D. I. (2006). A systematic review of statistical power in software engineering experiments. *Information and Software Technology*, *48*(8), 745–755.

Easterbrook, S., Singer, J., Storey, M.-A. & Damian, D. (2008). Selecting empirical methods for software engineering research. In *Guide to advanced empirical software engineering* (pp. 285–311). Springer.

Ebert, C. (1998). Putting requirement management into praxis: dealing with nonfunctional requirements. *Information and Software technology*, *40*(3), 175–185.

Ebrahimi, K., Jones, G. F. & Fleischer, A. S. (2014). A review of data center cooling technology, operating conditions and the corresponding low-grade waste heat recovery opportunities. *Renewable and Sustainable Energy Reviews*, *31*, 622–638.

Eckert, E. R. G. & Drake Jr, R. M. (1987). Analysis of heat and mass transfer.

Elashoff, J. D. (1981). Down with multiple t-tests! *Gastroenterology*, *80*(3), 615–619.

Ellison, G. N. (1984). *Thermal computations for electronic equipment*. Van Nostrand Reinhold Company.

Ellsworth, M. J., Goth, G. F., Zoodsma, R. J., Arvelo, A., Campbell, L. A. & Anderl, W. J. (2012). An overview of the ibm power 775 supercomputer water cooling system. *Journal of Electronic Packaging*, *134*(2), 020906.

Eustis, S. (2009, November 17). *Apparatus and method for using distributed servers as mainframe class computers.* Google Patents. (US Patent App. 12/620,579)

Facebook. (2010). New cooling strategies for greater data center energy efficiency. Retrieved 11 December 2016, from `https://www.facebook.com/notes/facebook-engineering/new-cooling-strategies-for-greater-data-center-energy-efficiency/448717123919/`

Fagiolini, A., Rocca, P., De Giorgi, S., Spina, E., Amodeo, G. & Amore, M. (2016). Clinical trial methodology to assess the efficacy/effectiveness of long-acting antipsychotics: randomized controlled trials vs naturalistic studies. *Psychiatry Research*.

Felter, W. M., Keller, T. W., Kistler, M. D., Lefurgy, C., Rajamani, K., Rajamony, R., ... Van Hensbergen, E. (2003). On the performance and use of dense servers. *IBM Journal of Research and Development*, *47*(5.6), 671–688.

Fernandez, K., Pyzdrowski, B., Schiller, D. W., Smith, M. B. et al. (2002). Understand the basics of centrifugal pump operation. *CEP Magazine, May*, *5256*.

Fozdar, F. M., Parkar, G. J. & Imberger, J. (1985). Matching temperature and conductivity sensor response characteristics. *Journal of physical oceanography*, *15*(11), 1557–1569.

France, R. & Rumpe, B. (2007). Model-driven development of complex software: A research roadmap. In *2007 future of software engineering* (pp. 37–54).

Gao, J. (2014). Machine learning applications for data center optimization.

Gawer, A. & Cusumano, M. A. (2002). *Platform leadership: How intel, microsoft, and cisco drive industry innovation*. Harvard Business School Press Boston.

Geet, O. (2014). Trends in data center design–ashrae leads the way to large energy savings. In *Nrel, ashrae conference, denver, june* (Vol. 24).

Gobor, A. (2016). What is pwm and how does it work? Retrieved 07 December 2016, from `https://www.ekwb.com/blog/what-is-pwm-and-how-does-it-work/`

Goh, H., Looi, M. & Kok, B. (2009). Comparison between direct-on-line, star-delta and auto-transformer induction motor starting method in terms of power quality. In *Proceedings of the international multiconference of engineers and computer scientists* (Vol. 2, pp. 18–20).

Golberg, A. & Robson, A. (1980). *Smalltalk: the language and its implentation.* Addison-Wesley.

Gomaa, H. (1989). Structuring criteria for real time system design. In *Proceedings of the 11th international conference on software engineering* (pp. 290–301).

Gomaa, H. (1993). *Software design methods for concurrent and real-time systems*. Addison-Wesley Longman Publishing Co., Inc.

Gomaa, H. (2001). Designing concurrent, distributed, and real-time applications with uml. In *Proceedings of the 23rd international conference on software engineering* (pp. 737–738).

Google. (2016a). Efficiency: How we do it. Retrieved 11 December 2016, from `https://www.google.com/about/datacenters/efficiency/internal/`

Google. (2016b). Solving for sustainability. Retrieved 07 December 2016, from `https://www.google.com/green/approach/`

Google. (2017a). Deepmind ai reduces google data centre cooling bill by 40%. Retrieved 23 January 2017, from `https://deepmind.com/blog/deepmind-ai-reduces-google-data-centre-cooling-bill-40/`

Google. (2017b). Google's latest data centre is cooled entirely with ocean water. Retrieved 23 January 2017, from `http://www.businessinsider.com.au/googles-latest-data-center-is-cooled-entirely-with-ocean-water-2011-5?r=US&IR=T`

Gousseau, P., Blocken, B. & Van Heijst, G. (2011). Cfd simulation of pollutant dispersion around isolated buildings: On the role of convective and turbulent mass fluxes in the prediction accuracy. *Journal of Hazardous Materials*, *194*, 422–434.

Greenberg, A., Hamilton, J., Maltz, D. A. & Patel, P. (2008). The cost of a cloud: research problems in data center networks. *ACM SIGCOMM computer communication review*, *39*(1), 68–73.

Halstead, M. H. (1977). Elements of software science: Elsevier north-holland. *New York*.

Hastings, R. J., Varghese, P. T., Jasso, T. P. & Leigh, K. B. (2002, October 22). *Modular rack-mount server arrangement.* Google Patents. (US Patent 6,469,899)

Hehenberger, P., Vogel-Heuser, B., Bradley, D., Eynard, B., Tomiyama, T. & Achiche, S.

(2016). Design, modelling, simulation and integration of cyber physical systems: Methods and applications. *Computers in Industry*, *82*, 273–289.

Intel. (1998). Introduction to thermal solutions. *Technical Notes by Intel Corp.*.

ISO, I. (2010). Ieee, systems and software engineering–vocabulary. *IEEE computer society, Piscataway, NJ*.

Jaeger, P. T., Lin, J. & Grimes, J. M. (2008). Cloud computing and information policy: Computing in a policy cloud? *Journal of Information Technology & Politics*, *5*(3), 269–283.

Jannou, C. (2015). *The internationalists: Globalization, the age of information and the developing world ascent*. Lulu Press, Inc.

Jedlitschka, A., Ciolkowski, M. & Pfahl, D. (2008). Reporting experiments in software engineering. In *Guide to advanced empirical software engineering* (pp. 201–228). Springer.

Jedlitschka, A. & Pfahl, D. (2005). Reporting guidelines for controlled experiments in software engineering. In *2005 international symposium on empirical software engineering, 2005.* (pp. 10–pp).

Jinkyun, C., Lim, T. & Kim, B. S. (2009). Measurements and predictions of the air distribution systems in high compute density (internet) data centers. *Energy and Buildings*, *41*(10), 1107–1115.

Jinkyun, C., Lim, T. & Kim, B. S. (2012). Viability of datacenter cooling systems for energy efficiency in temperate or subtropical regions: Case study. *Energy and buildings*, *55*, 189–197.

Kakac, S., Liu, H. & Pramuanjaroenkij, A. (2012). *Heat exchangers: selection, rating, and thermal design*. CRC press.

Kang, S., Miller, D. & Cennamo, J. (2007). Closed loop liquid cooling for high performance computer systems. In *Asme 2007 interpack conference collocated with the asme/jsme 2007 thermal engineering heat transfer summer conference* (pp. 509–515).

Kaplan, B. & Duchon, D. (1988). Combining qualitative and quantitative methods in information systems research: a case study. *MIS quarterly*, 571–586.

Karnouskos, S. (2011). Stuxnet worm impact on industrial cyber-physical system security. In *Iecon 2011-37th annual conference on ieee industrial electronics society* (pp. 4490–4494).

Kearns, M. J. (2010). Designing a digital future: Federally funded research and development in networking and information technology.

Keele, S. (2007). Guidelines for performing systematic literature reviews in software engineering. In *Technical report, ver. 2.3 ebse technical report. ebse.*

Kelle, U. & Erzberger, C. (2004). Qualitative and quantitative methods. *A companion to qualitative research. London. Sage*, 172–177.

Kitchenham, B., Al-Khilidar, H., Babar, M. A., Berry, M., Cox, K., Keung, J., … Zhu, L. (2008). Evaluating guidelines for reporting empirical software engineering studies. *Empirical Software Engineering*, *13*(1), 97–121.

Kitchenham, B., Brereton, O. P., Budgen, D., Turner, M., Bailey, J. & Linkman, S. (2009). Systematic literature reviews in software engineering–a systematic literature review. *Information and software technology*, *51*(1), 7–15.

Kitchenham, B. A., Dyba, T. & Jorgensen, M. (2004). Evidence-based software engineering. In *Proceedings of the 26th international conference on software engineering* (pp. 273–281).

Klein, L., Singh, P., Schappert, M., Griffel, M. & Hamann, H. (2011). Corrosion management for data centers. In *Semiconductor thermal measurement and management symposium (semi-therm), 2011 27th annual ieee* (pp. 21–26).

Koomey, J. (2011). Growth in data center electricity use 2005 to 2010. *A report by Analytical Press, completed at the request of The New York Times*, 9.

Koomey, J., Berard, S., Sanchez, M. & Wong, H. (2011). Implications of historical trends in the electrical efficiency of computing. *IEEE Annals of the History of Computing*, *33*(3), 46–54.

Kossiakoff, A., Sweet, W. N., Seymour, S. & Biemer, S. M. (2011). *Systems engineering principles and practice* (Vol. 83). John Wiley & Sons.

Kotonya, G. & Sommerville, I. (1998). *Requirements engineering: processes and techniques*. Wiley Publishing.

Kreeley, B. & Coulton, S. (n.d.). Increasing the lifespan and reliability of electrical components. *Kooltronic, Inc*.

Kumari, K. (2014). Software testing for finding errors. *International Journal Of Scientific Research And Education*, *2*(02).

Lachin, J. M. (1981). Introduction to sample size determination and power analysis for clinical trials. *Controlled clinical trials*, *2*(2), 93–113.

Lawrence, I. & Lin, K. (1989). A concordance correlation coefficient to evaluate reproducibility. *Biometrics*, 255–268.

Lee, E. A. (2008). Cyber physical systems: Design challenges. In *Object oriented real-time distributed computing (isorc), 2008 11th ieee international symposium on* (pp. 363–369).

Lin, L. & Ponnappan, R. (2003). Heat transfer characteristics of spray cooling in a closed loop. *International Journal of Heat and Mass Transfer*, *46*(20), 3737–3746.

Littlewood, B. & Verrall, J. (1973). A bayesian reliability growth model for computer software. *Applied statistics*, 332–346.

Madhavi, R. (2014). Liquid cooling vs. air cooling evaluation in the maui high performance computing center. *U.S. Department of Energy Federal Energy Management Program.*.

Mahapatra, S. & Yuan, X. (2010). Load balancing mechanisms in data center networks. In *the 7th int. conf. & expo on emerging technologies for a smarter world (cewit).*

Majidi, K. (2004). Numerical study of unsteady flow in a centrifugal pump. In *Asme turbo expo 2004: Power for land, sea, and air* (pp. 805–814).

Marcinichen, J. B., Wu, D., Paredes, S., Thome, J. R. & Michel, B. (2014). Dynamic flow control and performance comparison of different concepts of two-phase on-chip cooling cycles. *Applied Energy*, *114*, 179–191.

Mason, R. L., Gunst, R. F. & Hess, J. L. (2003). *Statistical design and analysis of experiments: with applications to engineering and science* (Vol. 474). John Wiley

& Sons.
Mathew Milnes. (n.d.). The mathematics of pumping water aecom design build civil, mechanical engineering. *AECOM*.
McCabe, T. J. (1976). A complexity measure. *IEEE Transactions on software Engineering*(4), 308–320.
McCartney, S. (1999). *Eniac: The triumphs and tragedies of the world's first computer*. Walker & Company.
Moogk, D. R. (2012). Minimum viable product and the importance of experimentation in technology startups. *Technology Innovation Management Review*, *2*(3), 23.
Morgan, D. L. (2007). Paradigms lost and pragmatism regained methodological implications of combining qualitative and quantitative methods. *Journal of mixed methods research*, *1*(1), 48–76.
Moss, D. L. (2011). Facility cooling failure: How much time do you have? *Technical Whitepaper by Dell Inc.*.
Murugesan, S. (2008). Harnessing green it: Principles and practices. *IT professional*, *10*(1).
Musilli, J. & Ellison, I. B. (2012). Facilities design for high-density data centers. *Intel White Paper*.
Neter, J., Wasserman, W. & Kutner, M. H. (1974). Applied linear statistical models. richard d. irwin. *Inc., Homewood, IL*, *842*.
Nguyen, J., Knight, R. W., Bhavnani, S. H., Harris, D. K., Fincher, S., Gess, J. & Sridhar, A. (2015). A numerical study of multiphase dielectric fluid immersion cooling of multichip modules including effects of nucleation site density and bubble departure diameter functions. In *Asme 2015 international mechanical engineering congress and exposition* (pp. V08AT10A032–V08AT10A032).
Oreskes, N. (2004). The scientific consensus on climate change. *Science*, *306*(5702), 1686–1686.
Pakbaznia, E. & Pedram, M. (2009). Minimizing data center cooling and server power costs. In *Proceedings of the 2009 acm/ieee international symposium on low power electronics and design* (pp. 145–150).
Parolini, L., Sinopoli, B., Krogh, B. H. & Wang, Z. (2012). A cyber–physical systems approach to data center modeling and control for energy efficiency. *Proceedings of the IEEE*, *100*(1), 254–268.
Parolini, L., Tolia, N., Sinopoli, B. & Krogh, B. H. (2010). A cyber-physical systems approach to energy management in data centers. In *Proceedings of the 1st acm/ieee international conference on cyber-physical systems* (pp. 168–177).
Patankar, S. V. (2010). Airflow and cooling in a data center. *Journal of Heat Transfer*, *132*(7), 073001.
Patel, C. D. & Shah, A. J. (2005). *Cost model for planning, development and operation of a data center.*
Patel, C. D., Sharma, R., Bash, C. E. & Beitelmal, A. (2002). Thermal considerations in cooling scale high compute density data centers. In *Thermal and thermomechanical phenomena in electronic systems, 2002. itherm 2002. the eighth intersociety conference on* (pp. 767–776).

Patterns, C. (n.d.). Design pattern catalog.

Patterson, M. K. & Fenwick, D. (2008). The state of datacenter cooling. *Intel Corporation White Paper. Available at http://download. intel. com/technology/eep/data-center-efficiency/stateof-date-center-cooling. pdf*.

Peffers, K., Tuunanen, T., Rothenberger, M. A. & Chatterjee, S. (2007). A design science research methodology for information systems research. *Journal of management information systems*, *24*(3), 45–77.

Poole, P. H., Sciortino, F., Essmann, U. & Stanley, H. E. (1992). Phase behaviour of metastable water. *Nature*, *360*(6402), 324–328.

Primo, J. (2012). Shell and tube heat exchangers basic calculations.

Reeves, C. & Brewer, J. (1980). Hypothesis testing and proof by contradiction: an analogy. *Teaching Statistics*, *2*(2), 57–59.

Rezania, A., Rosendahl, L. & Andreasen, S. J. (2012). Experimental investigation of thermoelectric power generation versus coolant pumping power in a microchannel heat sink. *International Communications in Heat and Mass Transfer*, *39*(8), 1054–1058.

Rice, W. R. (1989). Analyzing tables of statistical tests. *Evolution*, *43*(1), 223–225.

Romadhon, R., Ali, M., Mahdzir, A. M. & Abakr, Y. A. (2009). Optimization of cooling systems in data centre by computational fluid dynamics model and simulation. In *Innovative technologies in intelligent systems and industrial applications, 2009. citisia 2009* (pp. 322–327).

Sanders, M. C., legal representative Vivian, M.-S., Nguyen, J. D., Pascarella, R. et al. (2006, November 21). *Redundant data and power infrastructure for modular server components in a rack.* Google Patents. (US Patent 7,138,733)

Scale, S. (2017). Hvac simulation software cae for heating, air conditioning and ventilation applications. Retrieved 07 December 2016, from `https://www.simscale.com/hvac/`

Schmidt. (2004). Challenges in electronic cooling—opportunities for enhanced thermal management techniques—microprocessor liquid cooled minichannel heat sink. *Heat Transfer Engineering*, *25*(3), 3–12.

Schmidt, D. C. (2006). Model-driven engineering. *COMPUTER-IEEE COMPUTER SOCIETY-*, *39*(2), 25.

Schmidt, R. (2005). Liquid cooling is back. *Electron. Cooling*, *11*(3), 34–38.

Schmidt, R., Chu, R. C., Ellsworth, M., Iyengar, M., Porter, D., Kamath, V. & Lehman, B. (2005). 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* (pp. 663–673).

Shames, I. H. & Shames, I. H. (1982). *Mechanics of fluids* (Vol. 2). McGraw-Hill New York.

Sharma, R. K., Bash, C. E., Patel, C. D., Friedrich, R. J. & Chase, J. S. (2005). Balance of power: Dynamic thermal management for internet data centers. *IEEE Internet Computing*, *9*(1), 42–49.

Sharp, H., Finkelstein, A. & Galal, G. (1999). Stakeholder identification in the requirements engineering process. In *Database and expert systems applications, 1999. proceedings. tenth international workshop on* (pp. 387–391).

Shi, J., Wan, J., Yan, H. & Suo, H. (2011). A survey of cyber-physical systems. In *Wireless communications and signal processing (wcsp), 2011 international conference on* (pp. 1–6).

Shrivastava, S. K., Sammakia, B., Schmidt, R., Iyengar, M. & Van Gilder, J. W. (2006). Experimental-numerical comparison for a high-density data center: Hot spot heat fluxes in excess of 500 w/ft2. In *Thermal and thermomechanical proceedings 10th intersociety conference on phenomena in electronics systems, 2006. itherm 2006.* (pp. 402–411).

Sickinger, D., Van Geet, O. & Ravenscroft, C. (2014). *Energy performance testing of asetek's rackcdu system at nrel's high performance computing data center* (Tech. Rep.). National Renewable Energy Laboratory (NREL), Golden, CO.

Simons, R. E. (1995). The evolution of ibm high performance cooling technology. In *Semiconductor thermal measurement and management symposium, 1995. semi-therm xi., eleventh annual ieee* (pp. 102–114).

Sims, R. E., Rogner, H.-H. & Gregory, K. (2003). Carbon emission and mitigation cost comparisons between fossil fuel, nuclear and renewable energy resources for electricity generation. *Energy policy*, *31*(13), 1315–1326.

Sinha, R. (2013). Model-based design of coordinated traffic controllers.

Sinha, R. & Kakodkar, A. (2006). Design and development of the ahwr—the indian thorium fuelled innovative nuclear reactor. *Nuclear Engineering and Design*, *236*(7), 683–700.

Sjøberg, D. I., Hannay, J. E., Hansen, O., Kampenes, V. B., Karahasanovic, A., Liborg, N.-K. & Rekdal, A. C. (2005). A survey of controlled experiments in software engineering. *IEEE transactions on software engineering*, *31*(9), 733–753.

Smiałek, M. & Nowakowski, W. (2015). *From requirements to java in a snap: Model-driven requirements engineering in practice*. Springer.

Steinke, M. E. & Kandlikar, S. G. (2006). Single-phase liquid friction factors in microchannels. *International Journal of Thermal Sciences*, *45*(11), 1073–1083.

Stocker, T. (2014). *Climate change 2013: the physical science basis: Working group i contribution to the fifth assessment report of the intergovernmental panel on climate change*. Cambridge University Press.

Subramanian, R. S. (2008). Heat transfer to or from a fluid flowing through a tube. *web2. clarkson. edu/.../subramanian/.../Convective% 20Heat% 20Transfe r*.

Sultan, N. (2010). Cloud computing for education: A new dawn? *International Journal of Information Management*, *30*(2), 109–116.

Treasure, T. & MacRae, K. (1998). Minimisation: the platinum standard for trials? *British Medical Journal*, *317*(7155), 362–364.

Tuma, P. E. (2010). The merits of open bath immersion cooling of datacom equipment. In *Semiconductor thermal measurement and management symposium, 2010. semi-therm 2010. 26th annual ieee* (pp. 123–131).

Uddin, M. & Rahman, A. A. (2011). Techniques to implement in green data centres

to achieve energy efficiency and reduce global warming effects. *International Journal of Global Warming*, *3*(4), 372–389.
Wai, R.-J., Lin, C.-Y., Duan, R.-Y. & Chang, Y.-R. (2008). High-efficiency power conversion system for kilowatt-level stand-alone generation unit with low input voltage. *IEEE Transactions on Industrial Electronics*, *55*(10), 3702–3714.
Wang, Q., Ping, P., Zhao, X., Chu, G., Sun, J. & Chen, C. (2012). Thermal runaway caused fire and explosion of lithium ion battery. *Journal of power sources*, *208*, 210–224.
Webb, D. (2008). The green data centre: Cooling. *Network World Canada*, *24*(8), N_A.
Webster, J. & Watson, R. T. (2002). Analyzing the past to prepare for the future: Writing a literature review. *MIS quarterly*, xiii–xxiii.
Wolverton, R. W. (1974). The cost of developing large-scale software. *IEEE Transactions on Computers*, *100*(6), 615–636.
Zeighami, R., Saunders, W. A., Coles, H. & Branton, S. (2014). Thermal performance modeling of hybrid liquid-air cooled servers. In *Thermal and thermomechanical phenomena in electronic systems (itherm), 2014 ieee intersociety conference on* (pp. 583–587).
Zhu, T., Ren, J.-R. & Li, M.-F. (2009). Influence of generalized and extended uncertainty principle on thermodynamics of frw universe. *Physics Letters B*, *674*(3), 204–209.

Appendix A

Additional information

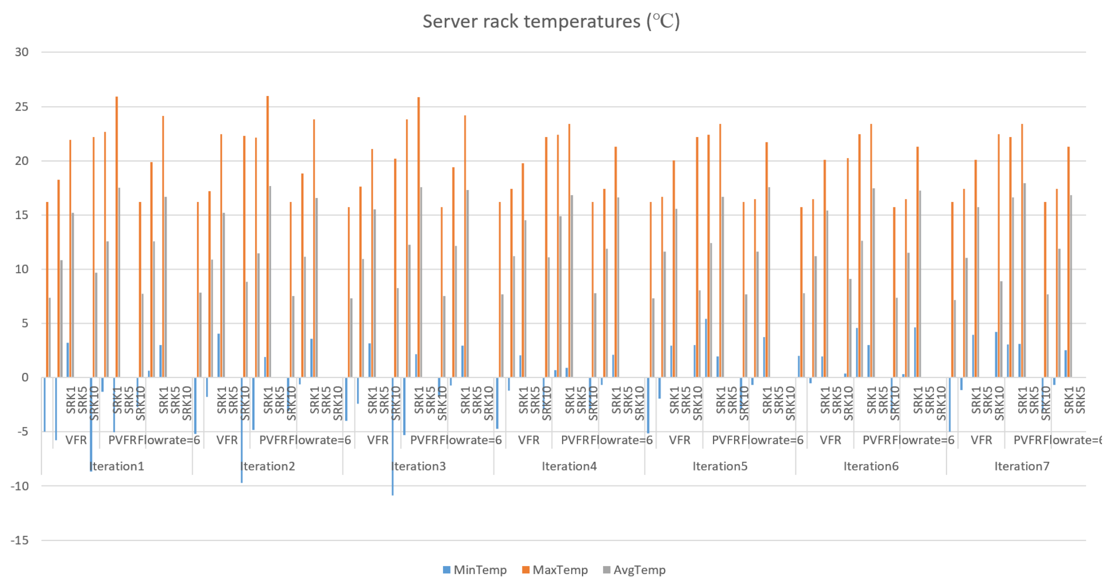

Figure A.1: Server rack temperatures in degrees Celsius for the three control strategies for each iteration while testing efficacy

A.1 URL/Link to DynaCool Code-base

`https://github.com/NidhiGowdra/DynaCool`

Please use the above URL/link to download the DynaCool system and experimental

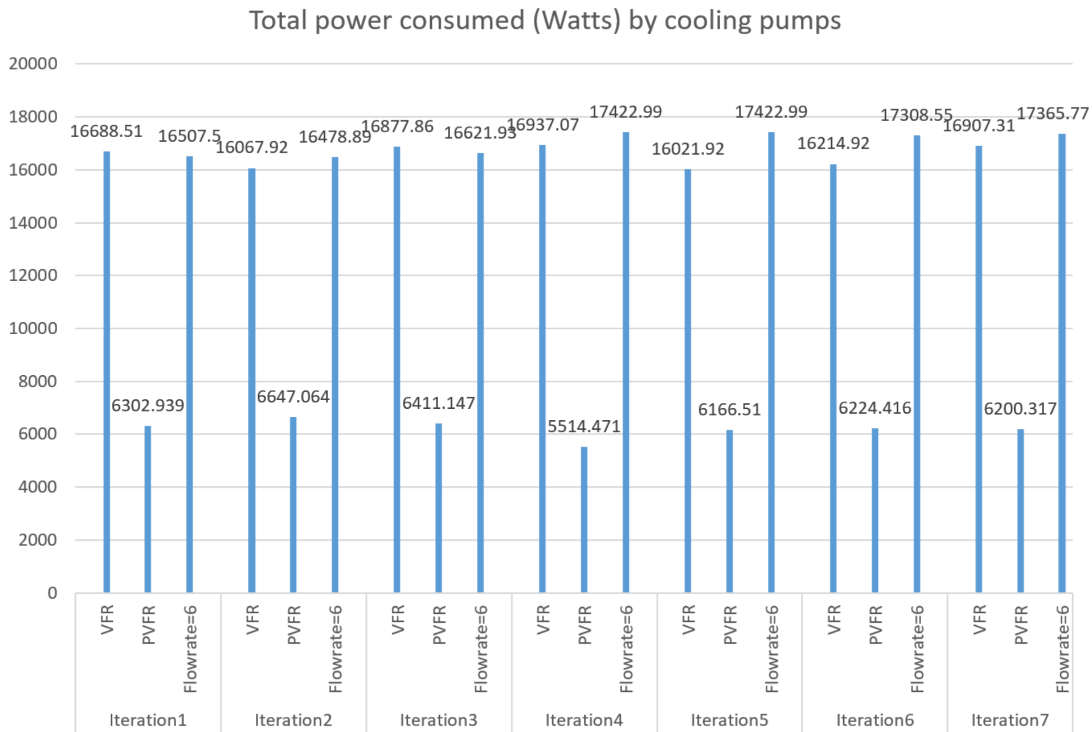

Figure A.2: Total power consumed by cooling pumps for the three control strategies for each iteration while testing efficacy

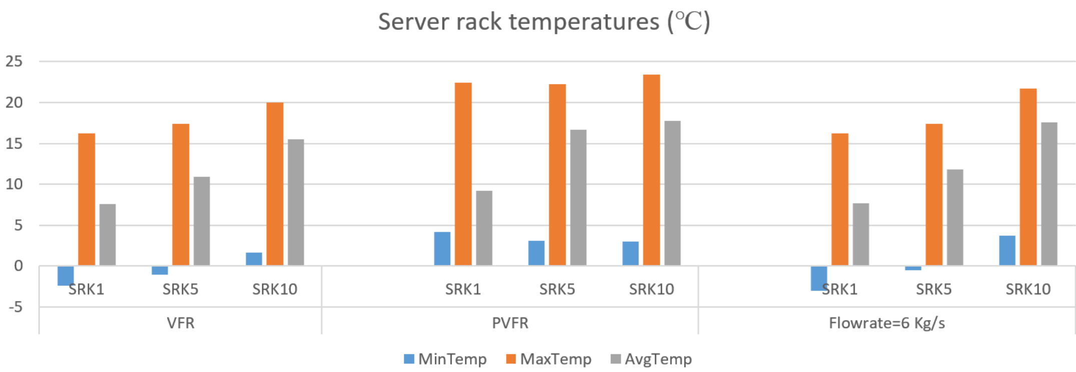

Figure A.3: Server rack temperatures in degrees Celsius for the three control strategies for a simulation time of 60 minutes

raw data. The link contains the DynaCool v2.1 software with the readme file, which specifies the instructions to operate the system along with some example files.

The link also contains raw experimental data in an archived (.rar) format, please use the free 'WinRAR' software to extract the files. The link to download the software is specified in the readme file.

A.2 Keywords used for literature search

The keywords used are listed below, in no particular order.

1. Liquid cooling
2. Data centre cooling
3. Next-generation data centres
4. Large scale data centres
5. Thermal guidelines
6. Computational Fluid Dynamics (CFD)
7. Data centre chillers
8. Cloud computing
9. Load balancing
10. Pump efficiency
11. Cooling efficiency
12. Thermodynamics in data centres
13. Cooling solutions for data centres
14. Green computing
15. Green solutions for data centres
16. Evaluation of cooling technologies in data centres

17. Dynamic cooling controllers

18. Dynamic liquid cooling

19. Variable flow rate cooling in data centres

20. Pulse width modulation

21. Operating thermal guidelines for data centres

22. Power density of server racks

23. Rack CDU

24. Google liquid cooling data centres