

EFFECT OF WARM WATER COOLING FOR AN ISOLATED
HYBRID LIQUID COOLED SERVER

by

ALEKHYA ADDAGATLA

Presented to the Faculty of the Graduate School of
The University of Texas at Arlington in Partial Fulfillment
of the Requirements
for the Degree of

MASTER OF SCIENCE IN MECHANICAL ENGINEERING

THE UNIVERSITY OF TEXAS AT ARLINGTON

December 2014

Copyright © by ALEKHYA ADDAGATLA 2014
All Rights Reserved

To my father Late Addagatla Suresh and mother Addagatla Anitha and my sister
Addagatla Hamsini.

ACKNOWLEDGEMENTS

I would like to thank my supervising professor Dr. Dereje Agonafer for constantly motivating and encouraging me, and also for his invaluable advice during the course of my study and research. I wish to thank my academic advisors Dr. Haji-Sheikh, Dr. Kent Lawrence, for taking time to serve in my thesis committee.

I thank Dr. Veerendra Mulay from Facebook Inc. for his technical guidance provided through my tenure at the University. I would like to thank John Edward Fernandes and Richard Eiland for their invaluable support, encouragement and for being my technical mentors in providing alternative solutions to benchmark the performance.

Finally, I would like to express my deep gratitude to my parents who have encouraged and sponsored for my graduate studies. I am extremely grateful to my mother for her sacrifice, encouragement and patience. I would also thank my friends, Divya and Ravi who helped me throughout my research.

November 6, 2014

ABSTRACT

EFFECT OF WARM WATER COOLING FOR AN ISOLATED HYBRID LIQUID COOLED SERVER

ALEKHYA ADDAGATLA, M.S.

The University of Texas at Arlington, 2014

Supervising Professor: Dr. Dereje Agonafer

Data center cooling strategies have changed overtime due to increases in power densities and ever increasing server cooling power requirement. Traditionally, most data centers adopt air cooling where the Computer Room Air Conditioning (CRAC) units pump volumes of chilled air to cool information technology (IT) equipment and dispose the heat from the servers to the ambient. With demand for energy-efficient and cost-effective data centers growing, alternative methods to both provide electrical power and cooling systems are the subject of active investigation. As data centers increase in size and current air-cooling systems reach their limitations, cooling technology will inevitably evolve towards air-liquid (hybrid) or liquid cooling systems.

Server level indirect cooling is achieved through a combination of warm water and recirculated air (hybrid cooling) in a 2OU (OpenU) web server. The work presented highlights the impact of various inlet supply water temperatures within the ASHRAE TC 9.9 liquid cooling classes (W4) ranging from 27.5°C to 45°C in terms of server power consumption, component temperatures and cooling power consumption.

TABLE OF CONTENTS

ACKNOWLEDGEMENTS	iv
ABSTRACT	v
LIST OF ILLUSTRATIONS	viii
LIST OF TABLES	x
Chapter	Page
1. INTRODUCTION	1
1.1 Data Center Efficiency	1
1.2 Need for liquid cooling	3
1.3 Motivation of the work	5
2. SYSTEM DESCRIPTION	7
2.1 Intel-based Open Compute Server	7
2.2 Current Scenario and Modification	10
3. COMPUTATIONAL FLUID DYNAMIC ANALYSIS	12
3.1 CFD Modeling in Data Centers	12
3.2 Fan Performance	13
3.3 Detailed CFD Server model	16
3.4 Final Duct Design	16
4. EXPERIMENTAL CHARACTERIZATION OF HYBRID LIQUID COOLED SERVER	20
4.1 Prototyped duct	20
4.2 Thermal Testing	23
4.3 Power Measurements	24

5. RESULTS AND DISCUSSION	27
5.1 Experimental Results of Hybrid Server	27
5.1.1 Results of CPU+MEM Tests	28
5.1.2 Results of 100% Tests	33
5.1.3 Results of 80% Tests	36
5.1.4 Results of 60% Tests	39
5.1.5 Results of 40% Tests	42
5.2 Summary of Results	45
6. CONCLUSION AND FUTURE WORK	46
6.1 Conclusion	46
6.2 Future Work	47
REFERENCES	48
BIOGRAPHICAL STATEMENT	52

LIST OF ILLUSTRATIONS

Figure		Page
1.1	Cooling in Data Center	2
1.2	Liquid-based Cooling at IBM	4
1.3	Layout of a traditional data center	5
1.4	Cooling system solution	6
2.1	System configuration	7
2.2	System at a glance	8
2.3	Server enclosed from ambient	9
2.4	Chassis cover with duct	10
2.5	ASHRAE T.c 9.9 Liquid Cooling Classes	11
3.1	CFD Modeling	12
3.2	CFD Analysis for duct design	13
3.3	Fan performance Curve	14
3.4	Fan performance Curve for SanACE 80	15
3.5	CFD showing original duct	17
3.6	CFD showing intermediate duct design	17
3.7	CFD showing final duct design	18
3.8	Streamline plot for Air Cooled Server with original duct	19
3.9	Streamline plot for Liquid Cooled Server with original duct	19
4.1	Improved duct for the server	21
4.2	Experimental test setup	21
4.3	Figure showing dry-cooler obstructions	22

4.4	Labview program to maintain inlet temperature	24
4.5	Labview program for measuring cooling power	24
4.6	Schematic of the Test setup	25
4.7	Complete test setup	26
5.1	Total Cooling Power for max CPU and Memory Workload	28
5.2	Plots for max CPU and Memory Workload	29
5.3	CPU Temperatures and Powers	30
5.4	Plots showing DIMM temperatures for DIMM 0 and DIMM 1	31
5.5	Plot for DIMM powers	32
5.6	Total Cooling Power for 100% CPU Workload	33
5.7	Plots for 100% CPU Workload	34
5.8	Plots showing CPU and DIMM Power	35
5.9	Plots showing DIMM temperatures for DIMM 0 and DIMM 1	35
5.10	Total Cooling Power for 80% CPU Workload	36
5.11	Plots for 80% CPU Workload	37
5.12	Plots showing CPU and DIMM Power	38
5.13	Plots showing DIMM temperatures for DIMM 0 and DIMM 1	38
5.14	Total Cooling Power for 60% CPU Workload	39
5.15	Plots for 60% CPU Workload	40
5.16	Plots showing CPU and DIMM Power	41
5.17	Plots showing DIMM temperatures for DIMM 0 and DIMM 1	41
5.18	Total Cooling Power for 40% CPU Workload	42
5.19	Plots for 60% CPU Workload	43
5.20	Plots showing CPU and DIMM Power	44
5.21	Plots showing DIMM temperatures for DIMM 0 and DIMM 1	44

LIST OF TABLES

Table		Page
3.1	Table showing Static Pressure and Differential Pressure	15
4.1	Table showing computational workloads	23
5.1	Consolidated data of the Server	45

CHAPTER 1

INTRODUCTION

1.1 Data Center Efficiency

Driven by the rapid increase in web activity through social media, web based businesses, media applications and services, data centers are quintessential in modern IT(Information Technology) infrastructure. With substantially growing energy demands for high performance computing architecture and associated equipment, data centers now represent a significant share of electricity in the U.S and the world [1]. The U.S Environmental Protection Agency (EPA), estimates that servers and data centers have consumed 1.5% of total US energy in 2006 [2]. Furthermore, in addition to these environmental deterioration and electricity bills, this increased power consumption may lead to system failures due to power over load or over heating.

Although the power distribution and cooling systems have reached its peak capacity, data centers continue to deploy high-density servers (e.g., blade servers) to cope with the ever-increasing workload. Power consumption in a data center is due a combination of: number of components installed on the server, CPU activity, memory, disk drives and also instructions being executed [3]. Thus, for optimized operations, design of a data center is given more importance, making data center research a prominent focus.

There has been a significant amount of research in improving the cooling efficiency of the data center including experimenting with the layout of equipment, dynamic monitoring and controls, localized air conditioning devices, right sizing [4]

and using virtual server technologies which allow the computational workloads to be better distributed.

Racked systems, which control the airflow within the system, can collectively require a different amount of volume air flow than the CRACs provide. However this leads to either high pressures and more re-circulation or inadequately provided air cooling in the room as seen in Fig 1.1. In the not so distant future, to match the high computational performance demands, high performance chips with heat fluxes of $100W/cm^2$ will increase heat loads at the server and facility level [5].




Figure 1.1. Cooling in Data Center.

A majority of the existing data centers use air-cooling systems, mainly chillers to maintain desired operating conditions. The heat generated by the processors is conducted to a heat sink and transferred to the cool air blowing into the server

through raised floor tiles. The cool air provided by the CRACs in turn cooled by chilled water, which is maintained at a sub-ambient temperature in order to provide sufficient heat transfer [6].

Most data centers have often preferred air-cooling due to its high reliability, wide range of computing, and lower initial and maintenance costs [7, 8]. However the amount of heat dissipated by the processors in air-cooled servers is too high to cool with the standard air-cooled heat sinks and fans.

By implementing hybrid or liquid cooling systems, cooling challenges can be met for improved efficiency and reducing energy and equipment costs [9]. Liquid cooling can serve higher heat densities and be more efficient than traditional air-cooling due to the higher heat capacities and low pumping power required for fluids.

1.2 Need for liquid cooling

Improved energy efficiencies will be realized when cooling systems allow the use of a warm water supply (77°F to 113°F rather than 60°F) to eliminate or reduce the need for mechanical cooling (chillers), described briefly in [10, 11]. Although the data center community has become conditioned to think of temperatures between 60 and 75 degrees as the proper climate for a server room, there are many ways to keep equipment running smoothly with cooling technologies featuring significantly higher temperatures.

These warmer water supply temperatures facilitate the pairing of liquid cooling with a water-side economizer (dry cooler), further increasing the potential of energy savings. The use of liquid cooling in such systems reduced the junction temperatures, CRAC units and thermal performance by 10% [12].



Figure 1.2. Liquid-based Cooling at IBM.

Data center cooling involves chillers that extract heat to the ambient with the assistance of a cooling tower, see Fig 1.3. However liquid cooling systems pose few trade-offs such as leaks bring a threat to the IT equipment, more infrastructure is required to contain the liquid and owing to these concerns, liquid cooling is more expensive when compared to air cooling [13]. Liquid cooling poses greater technical and budgetary challenges than air cooling but its effectiveness makes it a necessity for high-density applications.

While several reports have discussed the use of hybrid air-liquid cooling by servers and in data centers [14, 15, 16, 17], there has been little discussion of air flow and ducting for the auxiliary components.

1.3 Motivation of the work

Data center power use consists of information technology (IT) loads (such as servers, disk drives, and network equipment) and infrastructure loads (cooling, fans, pumps, lighting, and uninterruptible power supplies or UPSs). Traditionally, IT equipment consumes half of the data center energy with cooling infrastructure consuming about 20%-40% of the data center energy [18].

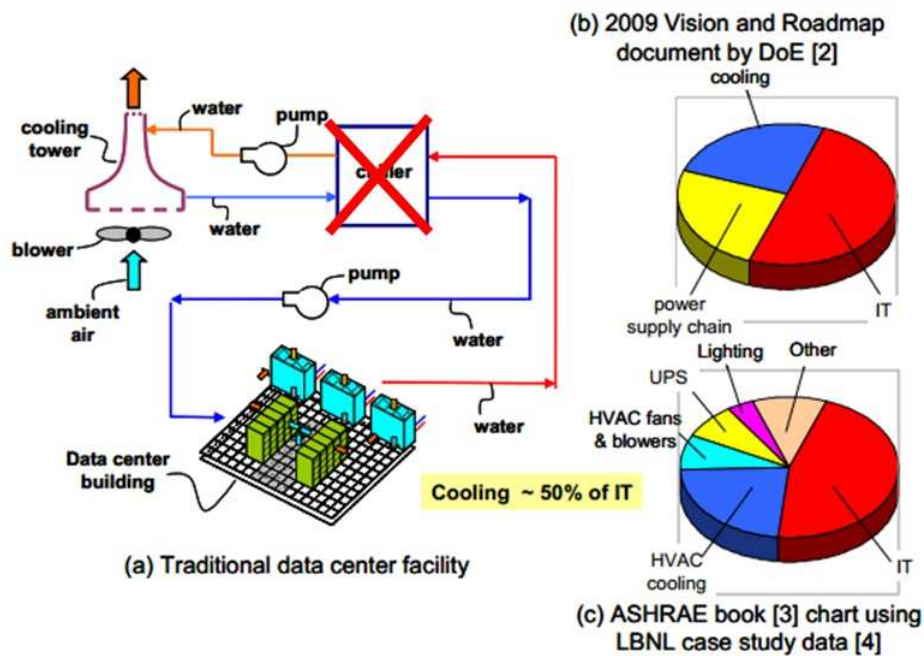


Figure 1.3. Layout of a traditional data center.

Most of this energy is consumed by the chiller plant, to provide chilled water to the data center and by CRAC and Computer Air Room Handlers (CRAH), to cool the computer room. Cooling system involves CRACs, chillers, cooling tower which contribute to a large consumption of energy in a data center. If the need for mechanical cooling (chillers) can be eliminated or greatly reduced without affecting

the operation, serviceability or reliability of IT equipment, significant amounts of energy can be saved [19, 20].

By implementing a cooling system shown in Fig 1.4, that utilizes liquid at temperatures above external ambient conditions (27.5°C to 45°C), to cool the majority of the heat generating components (CPUs in this case) in combination with traditional air cooling for remaining heat generating components such as Hard-disk drives (HDD), Platform Controller Hub (PCH) and other auxiliary components reduces the energy required to cool the electrical equipment, significantly reducing or even eliminating the need for chillers.

The specific heat carrying capacity of the liquid is four times that of air, resulting in lower mass flow rate, overall power draw and lower temperature rise. Thus, this can be advantageous to the existing data center infrastructure in presenting an energy efficient solution and potentially lowering system noise as a result of slower fan speeds. Further energy consumption could be realized by allowing warmer inlet air temperatures since most components require case temperatures of 70°C or higher with a ΔT of 20°C between the case and the liquid [19].




Figure 1.4. Cooling system solution.

CHAPTER 2

SYSTEM DESCRIPTION

2.1 Intel-based Open Compute Server

The server chassis with a form factor 2OU (88.9 mm) incorporates Intel Xeon 2570 dual CPUs, each dissipating a thermal design power (TDP) of 115W; specifications of the processor available in [21]. The server under study is referred to a CPU dominated server since the CPUs contribute to a majority of the thermal loads, refer to [22] for a detailed description of the server. In this setup, microprocessors conduct heat from the die through thermal interface materials (TIMs) to liquid cooled cold-plates, to enhance the heat transfer by replacing standard air cooled heat sinks shown in Fig 2.1.




Figure 2.1. System configuration.

The chassis is populated with 14 Dual-In-Line-Memory Module (DIMM) cards of 8 GB and dissipating a TDP of 3.9 W each approximately. The air flow inside the server enclosure was provided by an 80mm x 80mm fan attached to a heat exchanger (radiator) to cool the remaining heat generating components within the system. The server chassis includes a tray cover (lid) with an air duct to direct air flow through a radiator coil shown in Fig 2.2.




Figure 2.2. System at a glance.

The cooling system for this server consists of closed liquid loop which accommodates for the CPUs and air cooling for other heat generating components like DIMM cards, HDD, PCH, and auxiliary components on the printed-circuit board.

The processors in the server are indirect-liquid cooled with the cooling water temperature as high as 113°F (45°C). Low-power pumps integrated with each CPU cold plate provide flow for the system through external heat exchanger (dry cooler), internal radiator, and cold plates. Liquid maintained from 27°C to 45°C runs through fixed factory sealed tubes . Heat is directly dissipated from the microprocessors into the cooling water that flows through the cold plates. This minimizes the need for mechanical cooling (chillers).

The liquid coolant from the dry cooler first enters the radiator to cool the recirculating air flow in the system. After flowing through the radiator, the partially heated water enters the cold plates and recirculates back to the dry cooler to move the heat from the server to the ambient. The hot water coming out of the server is subsequently cooled to an ambient temperature of 25°C as it passes through a dry cooler. This server cooling design accommodates for both, air-cooled and liquid-cooled devices at the server level.

The key features of this server are, heat dissipated by the components is exhausted by the coolant and the server is secluded from the outside air environment, thus minimizing the risk of gaseous and component contamination.



Figure 2.3. Server enclosed from ambient.

2.2 Current Scenario and Modification

Major heat generating components being CPUs are cooled by liquid in the cold plates, while secondary components cooled by air. This allows hybrid liquid cooling to be compatible with the IT equipment that is air cooled as well as their liquid cooled counterparts. As mentioned earlier in the motivation (Section 1.3), this work aims to modify the system to a 100% liquid cooled server to reduce or eliminate the need for mechanical cooling (chillers). The modification in this study is to replace the primary duct to ensure adequate flow through the system in order to cool the auxiliary components.

The purpose of the duct shown in Fig 2.4 is to direct air flow through DIMMs placed to the left of the cold plates, pass over the Platform Controller Hub (PCH) and hard disk drive (HDD) and finally recirculate back through the other set of DIMMs to the right. Since the duct was designed for an air cooled server, the primary duct failed to accomplish this trend. It creates a low resistance across the chassis fan by recirculating the air.



Figure 2.4. Chassis cover with duct.

The scope of this study is to examine limitations of warm water cooling for this server within the American Society of Heating, Refrigerating and Air Conditioning Engineers (ASHRAE) TC 9.9 liquid cooling classes. In 2011, ASHRAE in collaboration with the IT equipment manufacturers proposed Thermal Guidelines for Liquid Cooled Data Processing Environments expanding their recommended environmental envelope for warmer liquid inlet temperatures.

Figure 2.5 shows liquid cooling classification for main and supplement cooling equipment. It can be observed that the chillers are associated with the main cooling equipment or supplement cooling equipment in the W1 to W3 classes. Whereas W4 envelope allows liquid temperatures ranging from 2°C (37.6°F) to 45°C (113°F) with the use of dry coolers or cooling tower. We focus on the W4 envelope for this system to facilitate warmer liquid temperatures ranging from 27.5°C to 45°C.

Liquid Cooling Classes	Main Cooling Equipment	Supplemental Cooling Equipment	Facility Supply water Temp (°F)
W1	Chiller/ Cooling Tower	Water-side economizer (dry cooler or cooling tower)	35.6 – 62.6
W2			35.6 – 80.6
W3	Cooling Tower	Chiller	35.6 – 89.6
W4	Water-side economizer (dry cooler or cooling tower)	N/A	37.6 – 113.0 (2°C – 45°C)
W5	Building Heating System	Cooling Tower	>113

Figure 2.5. ASHRAE T.c 9.9 Liquid Cooling Classes.

CHAPTER 3

COMPUTATIONAL FLUID DYNAMIC ANALYSIS

3.1 CFD Modeling in Data Centers

Computational Fluid Dynamics, more often abbreviated as CFD, is a branch of fluid dynamics tool that uses Finite Difference Method (FDM), Finite Element Method(FEM) and Finite Variable Method(FVM) to solve, analyze and predict complex fluid flows. Computers are used to simulate the physics and predict time dependent results given with initial boundary conditions.

However, it seems logical to use CFD as robust and dynamic tool to model data center environments as it gained immense popularity to model cooling effectiveness within racks and aisles. It provides a 3-D analysis of hot and cold air movement within the data center and identify hot spots or regions that require additional cooling or areas that are excessively cooled. CFD allows the customer to predict cooling power required for his application.




Figure 3.1. CFD Modeling.

Although CFD is used to analyze and design thermal air flow in data centers, there is indecisiveness with respect to accuracy of thermal predictions for large data centers. CFD modeling uses computer calculations to simulate the interactions of liquids and gases with various surfaces in a defined space [23].

In order to modify the system to a complete liquid cooled server, the initial step is to create a duct to cool DIMMs, PCH, HDD and other auxiliary heat generating components . Using 6SigmaET software code [24] by Future Facilities, candidate duct designs were simulated to best direct air flow through the heat generating components and recirculate through the DIMM cards. Figure 3.2 shows the air flow path through the final duct design.




Figure 3.2. CFD Analysis for duct design.

3.2 Fan Performance

Fans deliver air at a definite flow rate and static pressure, however the commonly used fans in electronic equipment are the axial fans, which deliver air parallel to the rotor axis and centrifugal blowers which deliver air perpendicular to the rotor axis. Generally axial fans are installed in servers since they have high flow rate capacity and work against low static pressure. Blowers are installed in laptops as it works against

higher static pressure and low flow rate capacity. Figure 3.3 shows an example for fan performance and fan efficiency curve.




Figure 3.3. Fan performance Curve.

Fan performance curve, also called as a fan curve, is a graphical representation of the air moving capacity of a fan. The radiator fan (San ACE 80) with 80mm x 80mm dimensions was tested in an air flow bench. Air flow bench measured the flow rate and static pressure simultaneously which determined the fan performance. The graph shown in Fig 3.4 is plotted with flow rate on abscissa (x - axis) and static pressure on ordinate (y - axis).

Table 3.1 shows static pressures obtained for differential pressures by increasing cfm. For a stall condition, all the nozzles are kept closed and the static pressure is measured. For free flow condition, the maximum flow rate is measured by adjusting blower blast gate to get static pressure frequency to zero.

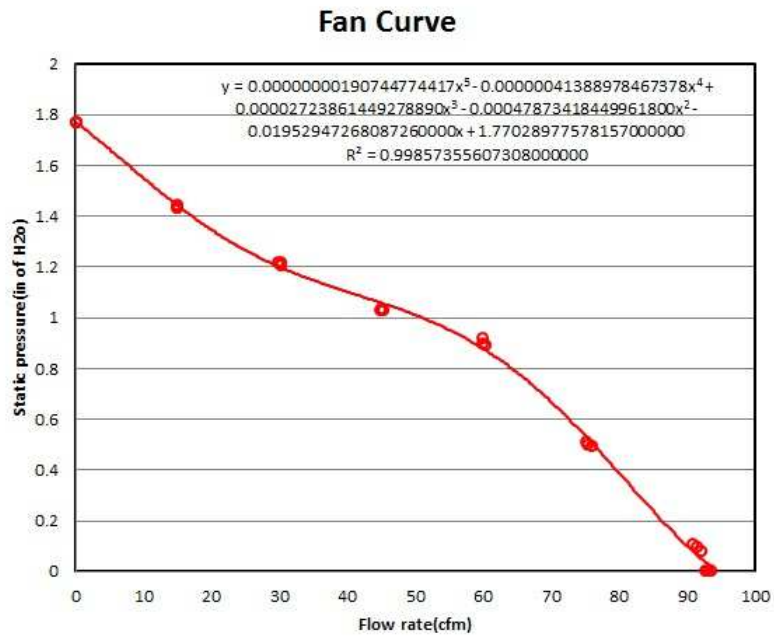


Figure 3.4. Fan performance Curve for SanACE 80.

Table 3.1. Table showing Static Pressure and Differential Pressure

CFM	Static Pressure	Differential Pressure
0	1.77	1.77
15	1.43	2.21
30	1.21	0.28
45	1.03	0.63
60	0.90	1.11
75	0.50	0.15
91	0.09	0.22
93	0.00	0.23

3.3 Detailed CFD Server model

A detailed CFD model is created using 6SigmaET [24], a commercially available CFD code which depicts the behavior to that of a data center. The server is modeled using all heat generating and flow impeding components including CPUs, cold plates, DIMMs, chipsets, server fan with radiator, and other auxiliary components such as mosfets, VRDs and chokes which are solved using K-E turbulence method. The model was completely isolated from the ambient to achieve a 100% liquid cooling of the server. CFD modeling is categorized into a flow model and a thermal(temperature) model. For this study, the model was solved for both flow and temperature.

For the software to solve for the given fan characteristics, fan performance curve and fan rpm's were given as input parameters for the fan. The radiator was tested in the air flow bench with surrounding duct to measure the resistance across he radiator. The resistance curve is plotted with flow rate on abscissa (x - axis) and static pressure on ordinate (y - axis) and given as input parameter for Z resistance. For a constant flow rate and fixed speed, the model was solved for flow and temperature solution.

3.4 Final Duct Design

CFD model helps in virtually prototyping duct designs to understand the thermal behavior of the server. Directing adequate flow to all the critical components such as DIMMs, HDD and PCH are important to air flow management. As discussed earlier, the duct on the chassis cover is designed for an air cooled server where majority of the cold air bypasses through the upstream region CPU 0 region, leading to an increase in CPU0 temperatures which is controlled by an inbuilt fan control algorithm. A set of duct designs are modeled and studied for the behavior of flow and temperature through system by the use of result phases available in the CFD code.

A CFM of 35 is prescribed for the flow inside the server. The termination factor was set to 0.1 to ensure accurate solution.

Several duct designs were modeled to achieve the desired flow pattern across the system. The duct design that created a recirculated flow as desired with lowest component temperatures and maximum flow rate was chosen as the final design. Original and intermediate duct designs are shown in the fig 3.5 and 3.6.




Figure 3.5. CFD showing original duct.




Figure 3.6. CFD showing intermediate duct design.




Figure 3.7. CFD showing final duct design.

Using the duct design shown in Fig 3.7, simulations were run for 3, 6, 12, 24 and 48 million cells by adjusting the grid size. Mesh sensitivity analysis is performed to ensure the solution is independent of the grid counts. It is observed that the model becomes grid independent at 24 million cells.

Figure 3.8 shows a streamline plot of the air cooled server model. The flow across the system is shown with velocity vectors in the form of arrows.




Figure 3.8. Streamline plot for Air Cooled Server with original duct.




Figure 3.9. Streamline plot for Liquid Cooled Server with original duct.

Figure 3,9 shows a streamline plot of the liquid cooled server model. It is shown from the figure that the fan forces air through first set of DIMMs, then passes through the PCH and HDD, recirculate back through the other set of DIMMs. Desired flow pattern was achieved.

CHAPTER 4

EXPERIMENTAL CHARACTERIZATION OF HYBRID LIQUID COOLED SERVER

4.1 Prototyped duct

CFD analysis plays a vital in understanding the flow patterns inside the server, however the CPU die temperatures and DIMM temperatures are inaccurate due to the fact that fan speeds and power inputs were fixed in the model, whereas fan speeds in a data center are controlled by fan control algorithm which is mostly triggered by the CPU die temperatures in the servers. The server in this study is slightly different from the air-cooled server, where the fan control algorithm gets triggered by the PCH temperatures. As the PCH temperature go beyond 72°C , the fan power increases thereby increasing the flow rate and cooling power. Thus it is mandatory to validate the duct design experimentally.


The goal of this implementation is to collect both thermal and power data to validate the water cooled server for high water inlet temperatures. T-type thermocouples placed at the radiator, DIMMs, HDD and the inlet hose of the dry cooler measured the temperature across these components. The remodeled duct design was prototyped using acrylic sheets of 3 mm size, see Fig 4.1. All the pieces were put together and then glued to the cover lid. The DIMMs and HDD are inserted to their respective slots, and server is sealed completely from the ambient.



Figure 4.1. Improved duct for the server.

To perform the study, thermocouple readings were collected through Data Acquisition (DAQ) and cooling power for the system was measured using a custom program (LabVIEW) that allowed control of water inlet temperature to the cold plates, by directing the dry-cooler fan speeds. In order to evaluate the cooling power, fan and the pump were powered externally and the power consumed was measured using LabVIEW [25].

The dry-cooler fans were controlled using an Arduino Uno microcontroller board and powered externally using a 12 VDC power supply, shown in Fig 4.2 (a) and (b) respectively.



(a) Dry Cooler Fans connected to arduino




(b) Power Supply (12VDC) to power dry cooler fans

Figure 4.2. Experimental test setup.

A foam board was used as resistance in order to reach the desired temperatures. Figure 4.3 (a) and (b) shows dry-cooler fans without and with obstructions respectively. For all tests, the radiators reject heat to the ambient laboratory environment maintained at $25^{\circ}\text{C} \pm 1.0^{\circ}\text{C}$. The DAQ and Arduino Uno are connected remotely to a computer workstation.



(a) Without Obstruction



(b) With Obstruction

Figure 4.3. Figure showing dry-cooler obstructions.

4.2 Thermal Testing

Thermal experiments are conducted on the server to evaluate for the CPU die temperatures, DIMM temperatures, cooling power consumption and total server power consumption. A USB data logger was used to measure the ambient laboratory temperature and relative humidity. The server was exercised computationally using a script with idle, 40%, 60%, 80%, 100% and at CPU+MEM power load, both the CPU and memory units were stressed simultaneously shown in table 4.1.

It should be noted that at this condition CPUs are stressed lower than 100% power level, i.e at 90% load and DIMM at 100% load. The test with 90% CPU and 100% memory power loads were conducted to provide continuous and steady heat dissipation from the processors and DIMMs. Since typical real-time workloads in data centers are not constant over long duration, the values presented here represent the worst-case scenario required to achieve steady-state condition.

Table 4.1. Table showing computational workloads

Computational Workloads		
NAME	CPU LOADING	MEMORY LOADING
Idle	Idle	Idle
40%	40%	
60%	60%	
80%	80%	
100%	100%	
CPU+MEM	90%	100%

4.3 Power Measurements

A LabVIEW program was written in order to maintain the inlet temperature of water, as shown in Fig 4.4. An upper limit and lower limit was specified for each water inlet temperature to be maintained. If the inlet temperature would exceed the upper limit, the dry cooler fans would ramp up in order to maintain the water temperature to be between the upper and lower limit specified.

In the same program power consumption of the server fan and pump were recorded which contribute to the cooling power consumption, see Fig 4.5. It should be noted that the cooling power consumption specifies the power consumed by both, the server fan and 3-wire pumps. Also the pumps were constant throughout the tests, resulting in constant flow across the cold plates.

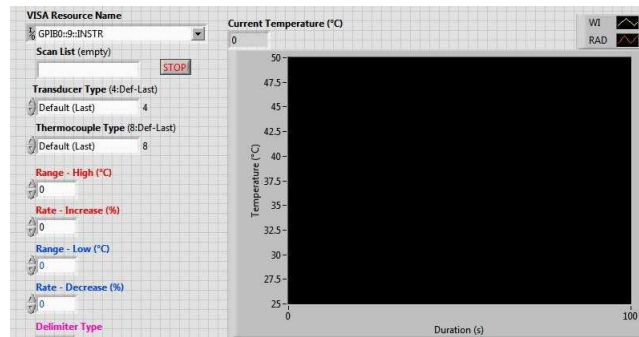


Figure 4.4. Labview program to maintain inlet temperature.

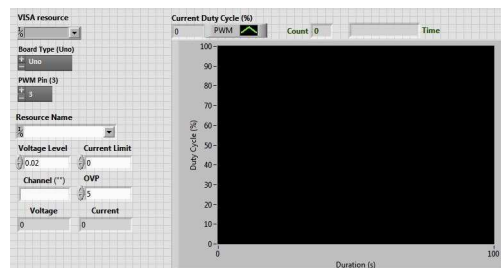


Figure 4.5. Labview program for measuring cooling power.

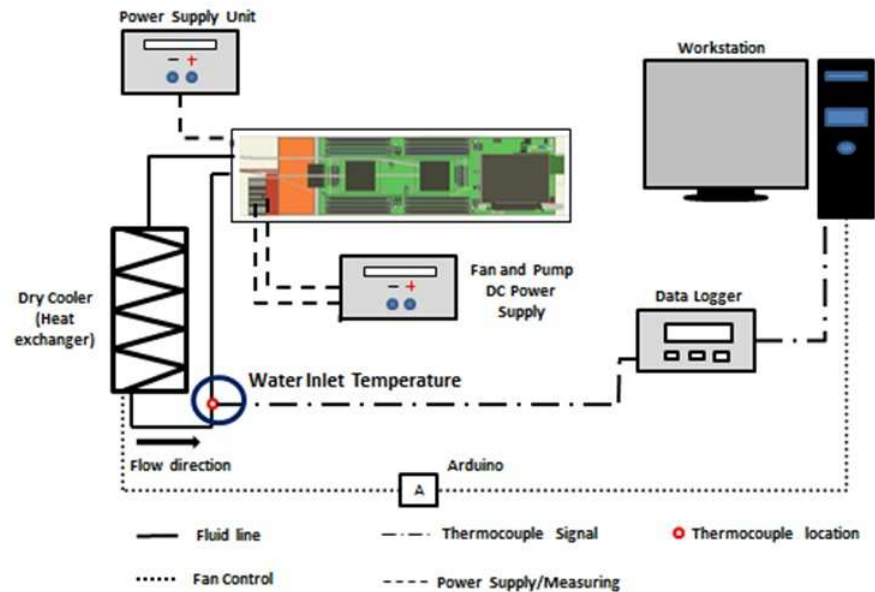


Figure 4.6. Schematic of the Test setup.

Figure 4.6 represents the schematic of the test setup showing all connections. The thermocouple from the dry cooler helped in reading the inlet temperature to the server.

The Power Thermal Utility (PTU) tools by Intel were used to measure, and load the CPU, memory units. Component information such as DIMM temperatures, powers, package powers, and total server power were recorded onto the workstation using ptumon and ipmitool. An internal diagnostic tool provided by the motherboard manufacturer reports data from DTS sensors in the processors, as well as, rpm readings from the internal server fan and pump.

Three test runs were conducted to ensure repeatability and consistency. An 18-hour test with three repeated trial runs of 6 hours each were conducted with different inlet water temperatures varying from 27.5°C to 45°C in increments of 2.5°C as shown in Fig 4.7.




Figure 4.7. Complete test setup.

Each power load were followed by a half hour idling period and one hour of precycle idling at the end of each trial run. Average of last ten minutes readings, at each power level is noted to account for accurate measurements after the fans have reached the stable condition. The results from maximum CPU and memory computational load were analyzed for drawing conclusions.

CHAPTER 5

RESULTS AND DISCUSSION

5.1 Experimental Results of Hybrid Server

This chapter discusses the results obtained by improving the duct design of the server and testing it experimentally. The hybrid-cooled server with prototyped duct was tested for water inlet temperatures ranging from 27.5°C to 45°C. Data recorded showed that all three runs occurred with water inlet temperatures within $\pm 0.5^\circ\text{C}$. As mentioned earlier, the pumps in the cold plates run at a fixed speed resulting in a constant flow rate.

All the data recorded by the ptumon, ipmitool and clock refresh was reduced to tabular columns for easy understanding. Data from all the inlet temperatures were consolidated into spreadsheets with idle, 40%, 60%, 80%, 100% and CPU+MEM computational loads. We consider CPU+MEM results for drawing conclusions since its considered as a worst case scenario. All the graphs are plotted with water inlet temperatures on x-axis and corresponding temperatures or powers on the y-axis.

IT power consumption includes only server power consumption and excludes the cooling power as the fan and pumps were powered and measured externally. The server power consumption includes the CPUs, HDD, PCH, DIMM cards and other auxiliary components on the board. The HDD consumes 7 W and DIMMs consumes 3.9 W each approximately.

5.1.1 Results of CPU+MEM Tests

Figure 5.1 shows data for total cooling power which is the sum of power drawn from the radiator fan and internal pumps. The graph is plotted with inlet temperatures obtained for three runs on x-axis and total cooling power on y-axis. The cooling power shows a piecewise linear relation with the increase in water inlet temperatures from 27.5°C to 45°C.

From the figure it shows that for water temperatures below 35°C cooling power increases gradually whereas for temperatures beyond 35°C it increases steeply until 45°C. Also with increase in water inlet temperatures from 27.5°C to 35°C and 35°C to 45°C, slopes of 0.075W/°C and 0.93 W/°C were observed.


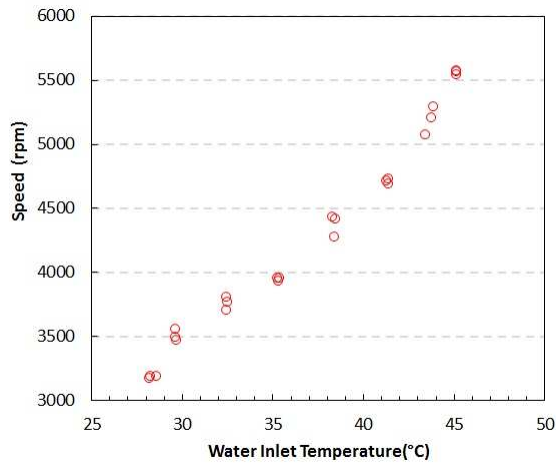
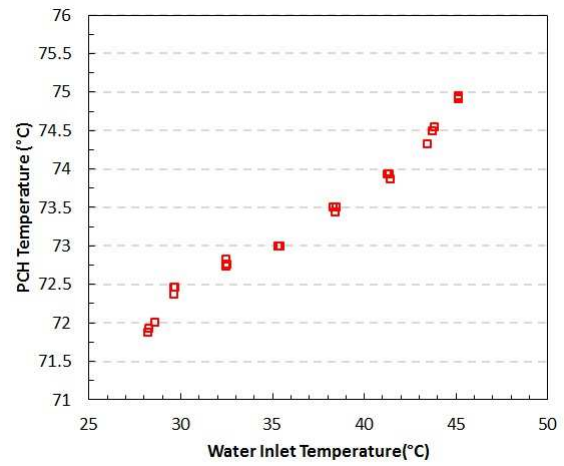



Figure 5.1. Total Cooling Power for max CPU and Memory Workload.



(a) Fan Speeds



(b) PCH Temperatures



(c) IT Power Consumption

Figure 5.2. Plots for max CPU and Memory Workload.

The fan control algorithm triggers the server fan to increase the flow rate beyond 35°C , thus resulting in higher cooling power. It is observed that when PCH temperatures rises beyond 71°C , radiator fan speeds increase with the rise in water inlet temperatures, shown in Fig 5.2 (b). The corresponding total cooling power is seen to be increasing with the increase in fan speeds.

The fan speeds operate at low speeds from 27.5°C to 35°C resulting in cooling power to increase gradually and operate at higher speeds beyond 35°C resulting in a sharp increase in cooling power as shown in Fig 5.1 (a).

The average IT power during this testing varied from 323 to 328 W, shown in Fig 5.2 (c). IT power consumption exhibits a linear relationship with rise in inlet water temperatures. The impact of leakage current at higher component temperatures was found to be minimal for this particular server.

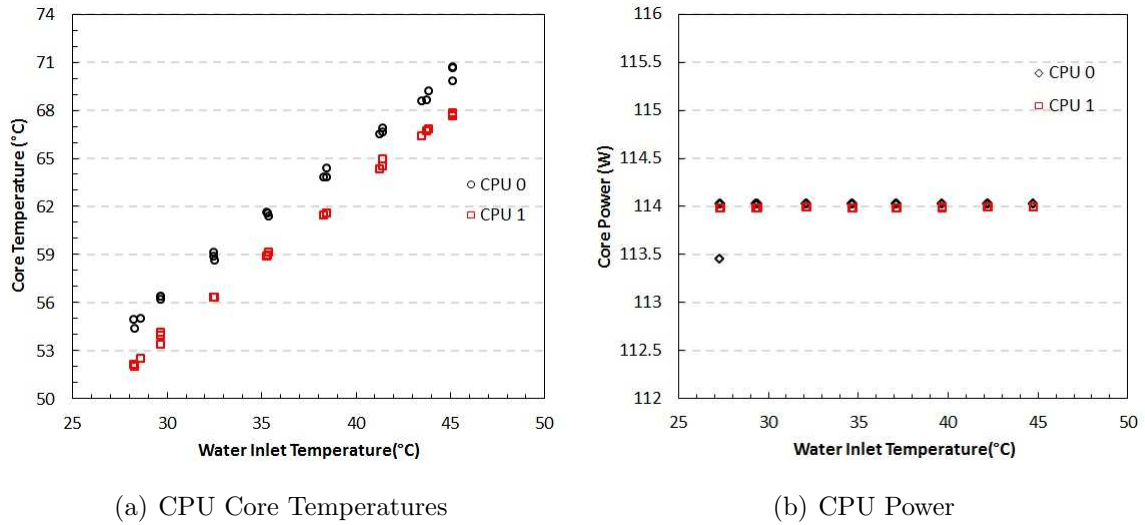


Figure 5.3. CPU Temperatures and Powers.

Fig 5.3 (a) shows the CPU temperatures as a linear function of water inlet temperatures. These higher temperatures were to be expected with rise in water temperatures; however they fall below the critical die temperatures. With every 2.5°C rise in inlet temperatures, the core temperatures increase in a similar trend due to the constant flow rate across the pumps. Although there is rise in water temperatures, a constant thermal performance is observed from the cold plates.

The power consumed by the CPUs throughout the CPU and memory test was approximately 114 W each as seen from the Fig 5.3 (b).

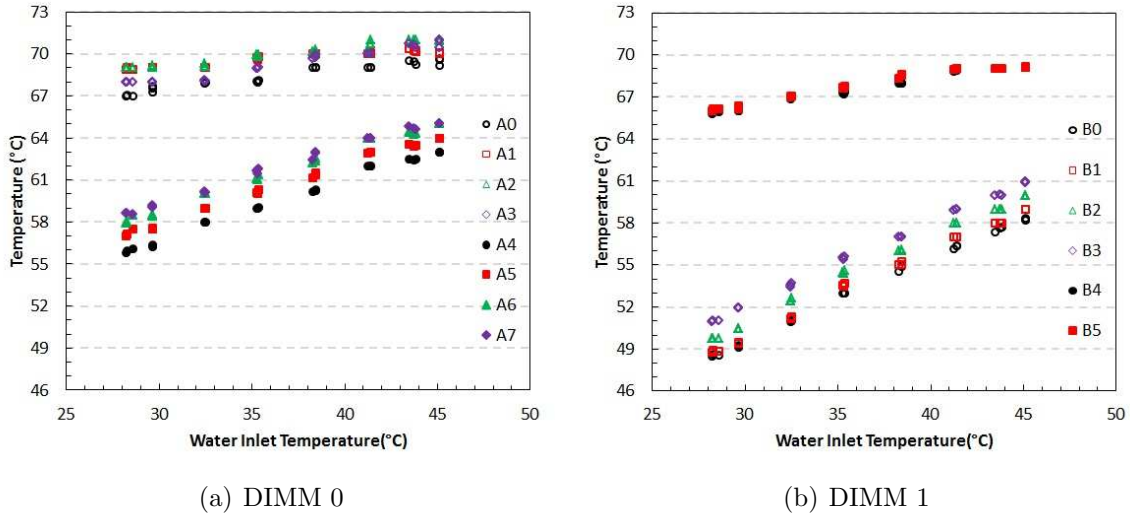



Figure 5.4. Plots showing DIMM temperatures for DIMM 0 and DIMM 1.



DIMM 0 and DIMM 1 recorded temperatures are shown in Fig 5.4 (a) and (b). Note that the DIMMs B0, B1, B2, B3, A4, A5, A6 and A7 closer to the internal radiator fan are cooled by cooler air, while the DIMMs A0, A1, A2, A3, B4 and B5 are away from the internal fan are cooled by warmer air due to the pre-heat from DIMMs

in left bank. The prototyped duct regulated the air flow through the DIMMs, HDD and other auxiliary components as expected and maintained component temperatures below the critical temperatures. The DIMM powers shown in Fig 5.5 were found to be decreasing for the rise in water inlet temperatures.

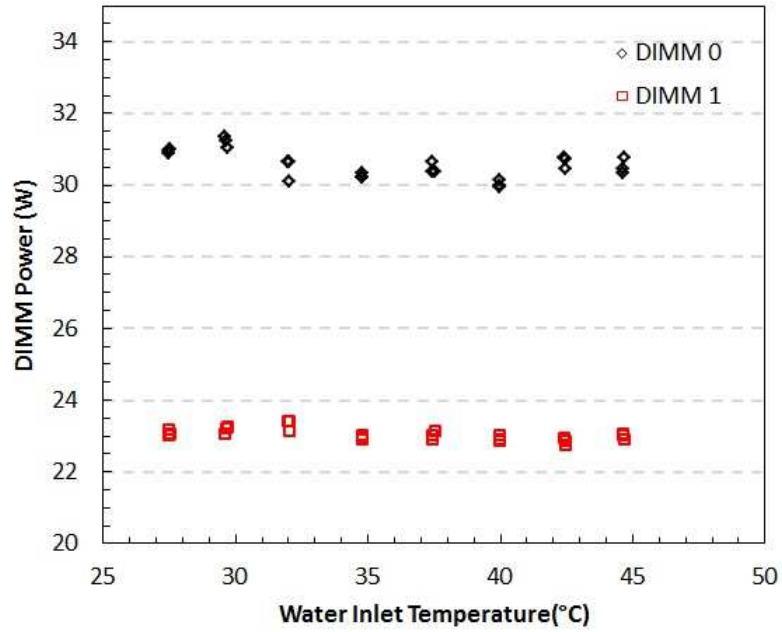


Figure 5.5. Plot for DIMM powers.

5.1.2 Results of 100% Tests

Figures show the effect of water inlet temperatures for a computational workload of 100% on the CPUs, however memory units were kept idle. The total cooling power for this test decreases initially and then gradually increases with the inlet water temperatures. A power of 4.24 W to 5.48 W is seen as the inlet temperatures rise from 27.5°C to 45°C. The corresponding IT power consumption varied from 288 W to 292 W which includes CPU power of 114 W each.







Figure 5.6. Total Cooling Power for 100% CPU Workload.




(a) Fan Speeds



(b) PCH Temperatures




(c) IT Power Consumption




(d) CPU Core Temperatures

Figure 5.7. Plots for 100% CPU Workload.




(a) CPU Power

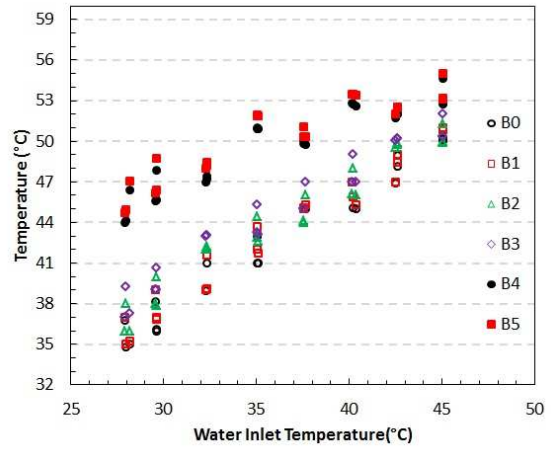


(b) DIMM Power

Figure 5.8. Plots showing CPU and DIMM Power.



(a) DIMM 0



(b) DIMM 1

Figure 5.9. Plots showing DIMM temperatures for DIMM 0 and DIMM 1.

5.1.3 Results of 80% Tests

Figures show the effect of water inlet temperatures for a computational workload of 80% on the CPUs, keeping memory units idle. The total cooling power for this test decreases initially and then gradually increases with the inlet water temperatures. A power of 4.23 W to 5.40 W is seen as the inlet temperatures rise from 27.5°C to 45°C and behaves similar to 100% CPU test. The corresponding IT power consumption varied from 276 W to 281 W which includes CPU power of 114 W each.







Figure 5.10. Total Cooling Power for 80% CPU Workload.




(a) Fan Speeds



(b) PCH Temperatures




(c) IT Power Consumption




(d) CPU Core Temperatures

Figure 5.11. Plots for 80% CPU Workload.




(a) CPU Power




(b) DIMM Power

Figure 5.12. Plots showing CPU and DIMM Power.



(a) DIMM 0



(b) DIMM 1

Figure 5.13. Plots showing DIMM temperatures for DIMM 0 and DIMM 1.

5.1.4 Results of 60% Tests

Figures show the effect of water inlet temperatures for a computational workload of 60% on the CPUs, keeping memory units idle. The total cooling power for this test decreases initially and then gradually increases with the inlet water temperatures. A power of 4.2 W to 4.72 W is seen as the inlet temperatures rise from 27.5°C to 45°C and behaves slightly different to 80% and 100% CPU tests. The corresponding IT power consumption varied from 232 W to 242 W which includes CPU power of 88 W - 91 W each.







Figure 5.14. Total Cooling Power for 60% CPU Workload.




(a) Fan Speeds



(b) PCH Temperatures




(c) IT Power Consumption

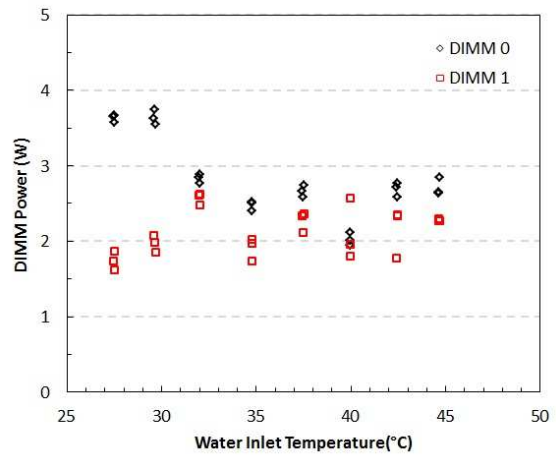


(d) CPU Core Temperatures

Figure 5.15. Plots for 60% CPU Workload.




(a) CPU Power

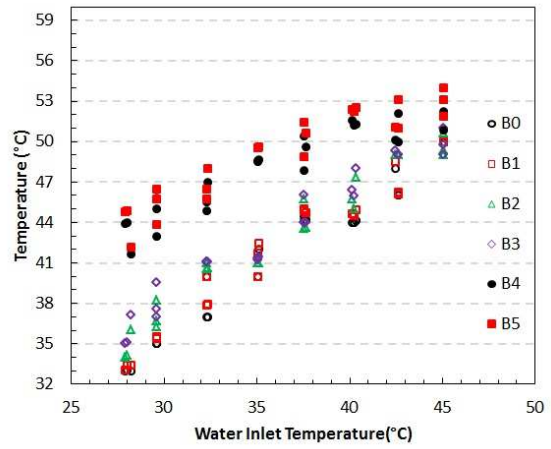


(b) DIMM Power

Figure 5.16. Plots showing CPU and DIMM Power.



(a) DIMM 0



(b) DIMM 1

Figure 5.17. Plots showing DIMM temperatures for DIMM 0 and DIMM 1.

5.1.5 Results of 40% Tests

Figures show the effect of water inlet temperatures for a computational workload of 40% on the CPUs, keeping memory units idle. The total cooling power for this test decreases with the inlet water temperatures. It decreases from 4.33 W to 4.13 W as the inlet temperatures rise from 27.5°C to 45°C. This is observed in every case irrespective of the fan sizes and power when the CPUs are stressed as low as 40%. The speeds are constant over the temperature range and the outcome is seen in the total cooling power. The corresponding IT power consumption varied from 182 W to 190 W which includes CPU power of 71 W to 77 W each.







Figure 5.18. Total Cooling Power for 40% CPU Workload.




(a) Fan Speeds



(b) PCH Temperatures




(c) IT Power Consumption




(d) CPU Core Temperatures

Figure 5.19. Plots for 60% CPU Workload.




(a) CPU Power




(b) DIMM Power

Figure 5.20. Plots showing CPU and DIMM Power.



(a) DIMM 0



(b) DIMM 1

Figure 5.21. Plots showing DIMM temperatures for DIMM 0 and DIMM 1.

5.2 Summary of Results

Table 5.1 represents the consolidated data of all the workloads showing CPU 0 and CPU 1 powers, DIMM 0 and DIMM 1 powers, total server power, total cooling power and fan speeds. For CPU+mem test conducted the total cooling power consumes 2.4% of the IT power.

Table 5.1. Consolidated data of the Server

Work-load(%)	CPU Power (W)	DIMM Power (W)	Server Power (W)	Total Cooling Power (W)	Fan Speeds (rpm)
CPU + Mem	C 0 : 113.5 - 114 C 1 : 114	D0 : 29.9-31.4 D1 : 22.8 - 23.4	323.6 - 328.5	4.9 - 7.8	3176 - 5576
100	C 0 : 114 C 1 : 114	D 0 : 1.1 - 3.6 D 1 : 1.6 - 2.4	287.9 - 292.1	4.2 - 5.5	2220 - 3896
80	C0: 113.5 - 114 C1: 113.4 - 114	D 0 : 2.0 - 3.9 D 1 : 1.5 - 2.8	276.3 - 281.7	4.2 - 5.4	2214 - 3576
60	C0 : 88.2 - 91.7 C1 : 85.5 - 88.8	D0 : 1.9 - 3.7 D1 : 1.6 - 2.6	231.5 - 241.9	4.2 - 4.7	2247 - 3180
40	C0 : 73.6 - 77.2 C1 : 71.9 - 74.8	D0 : 2.2 - 3.8 D1 : 1.6 - 2.6	181.7 - 190.3	4.1 - 4.3	2232 - 2362

CHAPTER 6

CONCLUSION AND FUTURE WORK

6.1 Conclusion

The main objective of this paper was to construct a 100% liquid-cooled system which could eliminate the need for mechanical cooling and this was achieved by experimental characterization of an enclosed hybrid-cooled server with a customized duct. CFD analysis helped in designing the customized duct which was tested experimentally to determine the impact of various water inlet temperatures on the power and thermal behavior of the system. 6SigmaET software enabled to understand the flow pattern for the duct design of the server.

Improved duct was prototyped using acrylic and glued together using acrylic glue. Experiments were conducted from idle, 40%, 60%, 80%, 100%, CPU and memory power levels with water inlet temperatures from 27.5°C to 45°C, however results from maximum CPU and memory setting was analyzed for drawing conclusions. From the results obtained, it is easy to notice there is ample incentive to operate the server for higher temperatures up to 45°C. This signifies that we could achieve to stay in W4 envelope recommended by ASHRAE.

For inlet water temperatures and server load dependent dry cooler fan control implemented, the average cooling power for all CPU+MEM tests was approximately 4.92 - 7.86 W with IT power of approximately 323 - 328 W leading to a cooling power fraction of 1.54 - 2.40% of IT power. The chip leakage throughout these test runs was insignificant, allowing more cooling power to yield better performance. The rpm of

the server fan increased predominantly based on the water inlet temperatures, CPU power levels, and PCH temperatures while the pump speed was held constant.

LabVIEW program helped in maintaining inlet temperatures from 27.5°C to 45°C and simultaneously measured cooling power which is a sum of server fan and cold plate pumps. Although the water inlet temperatures was raised to 45°C, the components were below their reliable temperatures. This determines that server-level indirect cooling enabled warm water cooling for supply water inlet temperatures as high as 45°C.

At CPU+MEM workload with a water inlet temperature of 45°C, the corresponding cooling power is <7.8 W, when compared to a typical data center running at lower server workloads. In summary, this experimental data provides validation of air duct design and thermal performance of the hybrid cooled server.

6.2 Future Work

The proposed system is an initial attempt to promote warm water cooling for an Open compute server. This study was conducted at system level, however this model can be extrapolated to facility level analysis and results for energy efficiency and performance can be examined. Also further enhancements to the system's performance may be gained by insulating the chassis interior surface to prevent heat loss to the ambient.

REFERENCES

- [1] J. G. Koomey, “Worldwide electricity used in data centers,” *Environmental Research Letters*, vol. 3, no. 3, p. 034008, 2008.
- [2] D. Patnaik, M. Marwah, R. Sharma, and N. Ramakrishnan, “Sustainable Operation and Management of Data Center Chillers Using Temporal Data Mining,” in *Proceedings of the 15th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, ser. KDD '09. New York, NY, USA: ACM, 2009, pp. 1305–1314. [Online]. Available: <http://doi.acm.org/10.1145/1557019.1557159>
- [3] “Rack and Power planning with HP Insight Control Power Management.”
- [4] R. Schmidt, “Liquid cooling is back,” *Electron. Cooling*, vol. 11, no. 3, pp. 34–38, 2005.
- [5] Y. J. Emad Samadiani and F. Mistree, “The Thermal Design of a Next Generation Data Center: A Conceptual Exposition,” *Journal of Electronic Packaging*, November 2008.
- [6] G. Meijer, “Cooling energy-hungry data centers,” *Science*, vol. 328, no. 5976, pp. 318–319, 2010.
- [7] S. Zimmermann, I. Meijer, M. K. Tiwari, S. Paredes, B. Michel, and D. Poulidakos, “Aquasar: A hot water cooled data center with direct energy reuse,” *Energy*, vol. 43, no. 1, pp. 237–245, 2012.
- [8] H. Coles, M. Ellsworth, and D. J. Martinez, “Hot for warm water cooling,” in *State of the Practice Reports*. ACM, 2011, p. 17.

- [9] S. Greenberg, E. Mills, B. Tschudi, P. Rumsey, and B. Myatt, “Best practices for data centers: Lessons learned from benchmarking 22 data centers,” *Proceedings of the ACEEE Summer Study on Energy Efficiency in Buildings in Asilomar, CA. ACEEE, August*, vol. 3, pp. 76–87, 2006.
- [10] M. Iyengar, M. David, P. Parida, V. Kamath, B. Kochuparambil, D. Graybill, M. Schultz, M. Gaynes, R. Simons, R. Schmidt, *et al.*, “Server liquid cooling with chiller-less data center design to enable significant energy savings,” in *Semiconductor Thermal Measurement and Management Symposium (SEMI-THERM), 2012 28th Annual IEEE*. IEEE, 2012, pp. 212–223.
- [11] T. ASHRAE, “Mission critical facilities, technology spaces, and electronic equipment, 2011, thermal guidelines for data processing environments, american society of heating,” *Refrigeration, and Air-Conditioning Engineers, Inc, Atlanta, GA*, 9.
- [12] J. Wei, “Hybrid cooling technology for large scale computing systems- from back to future,” in *Proceedings of ASME InterPACK 2011*.
- [13] M. Ohadi, S. Dessiatoun, K. Choo, M. Pecht, and J. Lawler, “A comparison analysis of air, liquid, and two-phase cooling of data centers,” in *Semiconductor Thermal Measurement and Management Symposium (SEMI-THERM), 2012 28th Annual IEEE*. IEEE, 2012, pp. 58–63.
- [14] M. Iyengar, M. David, P. Parida, V. Kamath, B. Kochuparambil, D. Graybill, M. Schultz, M. Gaynes, R. Simons, R. Schmidt, *et al.*, “Extreme energy efficiency using water cooled servers inside a chiller-less data center,” in *Thermal and Thermomechanical Phenomena in Electronic Systems (ITherm), 2012 13th IEEE Intersociety Conference on*. IEEE, 2012, pp. 137–149.
- [15] M. David, M. Iyengar, P. Parida, R. Simons, M. Schultz, M. Gaynes, R. Schmidt, and T. Chainer, “Experimental characterization of an energy efficient chiller-

- less data center test facility with warm water cooled servers,” in *Semiconductor Thermal Measurement and Management Symposium (SEMI-THERM), 2012 28th Annual IEEE*. IEEE, 2012, pp. 232–237.
- [16] P. R. Parida, M. David, M. Iyengar, M. Schultz, M. Gaynes, V. Kamath, B. Kochuparambil, and T. Chainer, “Experimental investigation of water cooled server microprocessors and memory devices in an energy efficient chiller-less data center,” in *Semiconductor Thermal Measurement and Management Symposium (SEMI-THERM), 2012 28th Annual IEEE*. IEEE, 2012, pp. 224–231.
- [17] M. David, M. Iyengar, P. Parida, R. Simons, M. Schultz, M. Gaynes, R. Schmidt, and T. Chainer, “Impact of operating conditions on a chiller-less data center test facility with liquid cooled servers,” in *Thermal and Thermomechanical Phenomena in Electronic Systems (ITherm), 2012 13th IEEE Intersociety Conference on*. IEEE, 2012, pp. 562–573.
- [18] R. Brown *et al.*, “Report to congress on server and data center energy efficiency:public law 109-431,” *Lawrence Berkeley National Laboratory*, 2008.
- [19] B. A. Rubenstein, R. Zeighami, R. Lankston, and E. Peterson, “Hybrid cooled data center using above ambient liquid cooling,” in *Thermal and Thermomechanical Phenomena in Electronic Systems (ITherm), 2010 12th IEEE Intersociety Conference on*. IEEE, 2010, pp. 1–10.
- [20] D. G. V. K. B. K. R. R. Milnes P. David, Robert Simons and R. Schmidt, “Cooling the cloud: Energy-efficient warm-water cooling of servers,” *Electronics Cooling:Dedicated to Thermal Management in the Electronics Industry*, 2013.
- [21] Intel, “<http://ark.intel.com/products/64595/>,” *Proceesor information*, 2012.
- [22] A. M. Harry Li, “<http://www.opencompute.org/assets/download/open-compute-project-intel-motherboard-v2.0.pdf>,” *Open Compute Project-Intel Motherboard Hardware v2.0*, 2011.

- [23] A. Almoli, A. Thompson, N. Kapur, J. Summers, H. Thompson, and G. Hannah, “Computational fluid dynamic investigation of liquid rack cooling in data centres,” *Applied energy*, vol. 89, no. 1, pp. 150–155, 2012.
- [24] FutureFacilities, “<http://www.air-think.com.tw/uploads/6/1/7/6/6176073/6sigmaet-whatsnew-r8.pdf>,” *6SigmaET software version 8*, 2011.
- [25] N. Instruments, “<http://www.ni.com/labview/>,” *Labview software*.

BIOGRAPHICAL STATEMENT

Alekhya Addagatla was born in Hyderabad, India, in 1990. She received her B.S. degree from JNTU University, India, in 2008. She qualified her Master of Science degree in Mechanical Engineering from University of Texas at Arlington in December 2014. During her masters program she has worked in thermal management of data centers and as a Teaching Assistant for Computer Aided Design lab. She has associated herself with industry collaborated research projects and studied various cooling techniques of electronic equipment. She has worked in experimental and CFD characterization of web cache servers, testing and validation of thermal solutions in cooling systems. She gained theoretical background knowledge of CFD and worked on data center specific CFD codes like 6SigmaET.