

# DEVELOPMENT OF AN EXPERIMENTAL APPARATUS FOR THE TESTING OF NOVEL COOLING SYSTEMS FOR PROCESSORS

Michael Hinton

Master of Engineering

Department of Mechanical Engineering

McGill University

Montréal, Québec

August 13, 2017

A thesis submitted to McGill University in partial fulfillment of the requirements of the  
degree of Master of Engineering

©Michael Hinton, 2017



## ACKNOWLEDGEMENTS

I would like to begin by thanking, most importantly, my supervisor, Professor Laurent Mydlarski. Through his continued support and guidance, I have grown as a scientific researcher and achieved my academic goals. I especially owe him a debt of gratitude for his role in bringing the present work to fruition.

This research project was conducted in collaboration with Hypertec, a computer hardware and services company in the Montréal area. Hypertec provided resource support in the form of funding, equipment and personnel. Thank you to Eliot Ahdoot for his role in initiating the project and seeing it through, and a special thank you to Fabrice, Jihad, Gilles, Florian, Noman, and Joseph for all of their technical and procurement advice and support.

Furthermore, I would like to voice my appreciation to the providers of financial support for this project. A McGill Engineering Undergraduate Student Masters Award (MEUSMA) was granted by the McGill Faculty of Engineering. Moreover, I was awarded an Alexander Graham Bell Canada Graduate Scholarship-Master's (CGS M) by the Natural Sciences and Engineering Research Council of Canada (NSERC), who also supplied an NSERC Engage grant. Lastly, Mitacs funded the project by means of a Mitacs Accelerate award.

I would be remiss if I did not also thank my friends and family for their endless support. Whether that support was encouragement or, at times, a much needed distraction, it was greatly appreciated in all of its forms. I would also like to thank Justin Riggio for his help in translating the abstract of this thesis.

## ABSTRACT

A computer system's performance is largely dependent on its ability to effectively remove the substantial heat generated by its processor, which is concentrated in a relatively small area of the system. While air cooling techniques remain sufficient for most personal computer needs, liquid cooling technologies, such as microchannels, have become commonplace in the most advanced computer systems. With the goal of further improving the performance of advanced processors, an experimental apparatus has been designed and built to investigate new cooling methods. The present work covers the design, construction and testing of this processor cooling experimental apparatus, as well as an initial prototype cooling device. The apparatus contains a fluid loop that can handle high-pressure coolant flow, and is suitable for both single-phase and two-phase cooling devices. Furthermore, the fluid loop is equipped with a positive displacement pump that provides very accurate flow rates, and multiple temperature and pressure sensors. A dedicated computer controller and custom software application operate the data acquisition and control systems that have been incorporated into the experimental apparatus. Moreover, the apparatus has been designed to conduct two different types of experiments. Initial tests use a heated surface that simulates the characteristics of a real processor, and is capable of providing data from power and temperature sensors that cannot be incorporated into a processor. The second set of tests use a real processor inside an advanced computer system, and they determine the processor performance using the new cooling system. Given the overall aim of improving processor cooling, multiple novel processor cooling device designs were conceived, and the most promising design was prototyped. The design uses a spray cooling method that directly targets the processor surface, thus eliminating thermal resistances that are inherent to most processor cooling devices. This thesis culminates with a presentation of the results of both types of tests using the prototype in the experimental apparatus.

## RÉSUMÉ

La performance d'un système d'ordinateur est largement dépendante de sa capacité d'enlever la chaleur provenant du processeur, qui se concentre sur une petite superficie, de manière efficace. Le refroidissement par air est une technique adéquate pour la majorité des ordinateurs personnels. Pour les systèmes les plus avancés qui requièrent un refroidissement plus important, les techniques à base de liquides, tels que les microcanaux, sont devenues populaires. Ayant l'objectif d'améliorer la performance des processeurs avancés, un appareil a été conçu et développé pour l'analyse de nouvelles méthodes de refroidissement. Ce travail comprend la conception, la construction et l'analyse de l'appareil, ainsi qu'un prototype initial. L'appareil comprend un circuit de fluide capable de supporter la haute pression du fluide de refroidissement, de phase unique ou biphasé. En outre, le circuit de fluide est équipé d'une pompe à déplacement positif, qui engendre des débits très précis, et de multiples capteurs de pression et de température. Un ordinateur dédié et un logiciel personnalisé, intégré dans l'appareil, gèrent l'acquisition des données et le système de contrôle. L'appareil est conçu pour conduire deux différents types d'expériences. Pour les tests initiaux, une surface chauffée simule les caractéristiques d'un processeur réel et fournit les données des capteurs de températures et de puissances électriques qui n'existent pas au sein d'un processeur. Le deuxième ensemble de tests utilise un processeur réel d'un système d'ordinateur avancé pour déterminer la performance de la nouvelle méthode de refroidissement. Ayant l'objectif d'améliorer la performance des processeurs avancés, plusieurs conceptions novatrices d'appareils ont été conçues et la conception la plus prometteuse a été réalisée en prototype. Ce modèle utilise une méthode de refroidissement par atomisation visant directement la surface du processeur afin d'éliminer la résistance thermique qui est présente dans la majorité des systèmes de refroidissement pour processeur. Ce mémoire culmine avec la présentation des résultats des deux types de tests utilisant le prototype dans l'appareil expérimental.

## TABLE OF CONTENTS

ACKNOWLEDGEMENTS . . . . .	iii
ABSTRACT . . . . .	iv
RÉSUMÉ . . . . .	v
LIST OF TABLES . . . . .	ix
LIST OF FIGURES . . . . .	x
NOMENCLATURE . . . . .	1
1 Introduction . . . . .	1
1.1 Background, Motivation and Overall Objectives . . . . .	1
1.2 Specific Objectives . . . . .	6
1.3 Literature Review . . . . .	7
1.3.1 Flat Surface Heat Transfer . . . . .	7
1.3.2 Cooling Methods . . . . .	8
1.3.3 Cooling Enhancement . . . . .	19
1.4 Overview of Thesis . . . . .	22
2 Experimental Apparatus . . . . .	23
2.1 Fluid Loop: Design and Construction . . . . .	23
2.1.1 Overview . . . . .	24
2.1.2 Pump . . . . .	26
2.1.3 Radiator and Fans . . . . .	28
2.1.4 Reservoirs . . . . .	28
2.1.5 Valves . . . . .	29
2.1.6 Sensors . . . . .	29
2.1.7 Tubing and Fittings . . . . .	31
2.1.8 Construction and Assembly . . . . .	33
2.2 Simulated CPU: Design and Construction . . . . .	34
2.2.1 Heated Surface . . . . .	34
2.2.2 Case and Insulation . . . . .	35
2.2.3 Temperature Sensors . . . . .	36
2.2.4 Electrical Power . . . . .	36
2.3 Real CPU: Computer System and Construction . . . . .	39
2.3.1 Computer System Specifications . . . . .	39

2.3.2	Computer System Software . . . . .	41
2.3.3	Stand Construction and Assembly . . . . .	44
2.4	Control and Data Acquisition . . . . .	44
2.4.1	Master Controller: Raspberry Pi . . . . .	45
2.4.2	Slave Controllers . . . . .	46
2.4.3	Inputs . . . . .	48
2.4.4	Outputs . . . . .	50
2.4.5	Application . . . . .	50
3	CPU Cooling Device . . . . .	55
3.1	Design Objectives and Requirements . . . . .	55
3.2	Conceptual Design Overview . . . . .	58
3.3	Prototype Detailed Design . . . . .	61
3.3.1	Coolant . . . . .	61
3.3.2	Nozzle . . . . .	63
3.3.3	Spray Chamber . . . . .	65
3.4	Assembly . . . . .	69
4	Methods . . . . .	71
4.1	Simulated CPU Experiments . . . . .	71
4.1.1	Setup . . . . .	71
4.1.2	Experiment . . . . .	72
4.1.3	Shutdown and Obtaining Data . . . . .	73
4.2	Real CPU Experiments . . . . .	74
4.2.1	Setup . . . . .	74
4.2.2	Experiment . . . . .	75
4.2.3	Shutdown and Obtaining Data . . . . .	77
5	Results . . . . .	78
5.1	Fluid Flow Rate and Pressure Experiments . . . . .	78
5.2	Simulated CPU Experiments . . . . .	80
5.2.1	Single-Phase Coolant: Distilled Water . . . . .	80
5.2.2	Two-Phase Coolant: FC-3284 . . . . .	83
5.3	Real CPU Experiments . . . . .	90
6	Conclusion . . . . .	99
6.1	Review of the Thesis . . . . .	99
6.2	Extensions of the Present Work . . . . .	101
	REFERENCES . . . . .	102
	APPENDIX A: Heat Transfer Rate Estimation . . . . .	108

APPENDIX B: Uncertainty Analysis . . . . .	110
B.1 Simulated CPU Surface Temperature . . . . .	111
B.2 Simulated CPU Case Temperature . . . . .	112
B.3 Spray Chamber Inlet and Outlet Temperatures . . . . .	113
B.4 Ambient Temperature . . . . .	115
B.5 Pressure . . . . .	115
B.6 Sensible Heat Transfer Rate . . . . .	116
B.7 Simulated CPU Power . . . . .	116
B.8 Estimated Heat Transfer Rate . . . . .	117
B.9 Vapour Quality . . . . .	118
B.10 Spray Efficiency . . . . .	118

## LIST OF TABLES

<u>Table</u>	<u>page</u>
2-1 CIARA’s ORION HF320-G3 server specifications. . . . .	40
2-2 Intel® Core™ i7-5960X Processor Extreme Edition (8 cores, 20 MB) specifications. . . . .	41
2-3 Input connections and their signal types. . . . .	49
2-4 Output connections and their signal types. . . . .	50
3-1 Intel® LGA-2011-3 socket vertical dimensions. . . . .	56
3-2 Relevant properties of distilled water and 3M’s Fluorinert™ Electronic Liquid FC-3284. . . . .	63
4-1 UEFI overclocking settings. . . . .	75
4-2 UEFI settings for specified CPU frequency. . . . .	77
5-1 Pass or fail results of the real CPU experiments for the Asetek, distilled water spray and FC-3284 spray cooling systems. . . . .	92
B-1 Summary of bias errors for the simulated CPU surface temperature measurements. . . . .	112
B-2 Summary of bias errors for the simulated CPU case temperature measurements. . . . .	113
B-3 Summary of bias errors for the ambient temperature measurements. . . . .	115
B-4 Summary of bias errors for the pressure transmitter measurements. . . . .	115

## LIST OF FIGURES

<u>Figure</u>	<u>page</u>
2-1 Experimental apparatus fluid loop diagram. . . . .	25
2-2 Qdos 30 Universal+ Peristaltic Metering Pump. . . . .	27
2-3 Experimental apparatus fluid loop - front view. . . . .	32
2-4 Experimental apparatus fluid loop - top view. . . . .	33
2-5 Simulated CPU with attached CPU cooling device. . . . .	35
2-6 Modified Kill A Watt Electricity Usage Monitor attached to custom breaker outlet. . . . .	38
2-7 Inside CIARA's ORION HF320-G3 server. . . . .	40
2-8 ASUS UEFI BIOS Utility graphical user interface. . . . .	42
2-9 OCCT and HWiNFO64 application windows. . . . .	43
2-10 Experimental apparatus with server in stand. . . . .	44
2-11 Window of the application's graphical user interface. . . . .	53
3-1 Intel® LGA-2011-3 socket exploded view. . . . .	57
3-2 Intel® LGA-2011-3 socket labeled view. . . . .	57
3-3 Intel® LGA-2011-3 socket vertical dimensions view. . . . .	58
3-4 Intel® Core™ i7 Processor IHS dimensions. . . . .	58
3-5 CPU spray cooling device conceptual design. . . . .	60
3-6 UniJet Full Cone Spray Nozzle exploded view. . . . .	64
3-7 Spray chamber wall - CAD - bottom isometric view. . . . .	67
3-8 Spray chamber lid - CAD - top isometric view. . . . .	69
3-9 Spray chamber assembled - top isometric view. . . . .	70
3-10 Spray chamber assembled - bottom view. . . . .	70
5-1 Spray nozzle pressure drop as a function of flow rate for distilled water. . . . .	79

5-2 Heat transfer rate as a function of surface temperature and flow rate for the simulated CPU using distilled water as the coolant. . . . . 82

5-3 Heat transfer rate as a function of surface temperature and flow rate for the simulated CPU using FC-3284 as the coolant. . . . . 85

5-4 Sensible heat transfer rate as a function of surface temperature and flow rate for the simulated CPU using FC-3284 as the coolant. . . . . 87

5-5 Vapour quality as a function of surface temperature and flow rate for the simulated CPU using FC-3284 as the coolant. . . . . 88

5-6 Spray efficiency as a function of surface temperature and flow rate for the simulated CPU using FC-3284 as the coolant. . . . . 90

5-7 Average CPU core temperature comparison for the real CPU experiments using the Asetek, distilled water spray, and FC-3284 spray cooling systems. 93

5-8 Maximum CPU core temperature comparison for the real CPU experiments using the Asetek, distilled water spray and FC-3284 spray cooling systems. 94

5-9 CPU power comparison for the real CPU experiments using the Asetek, distilled water spray, and FC-3284 spray cooling systems. . . . . 95

5-10 Sensible heat transfer rate comparison for the real CPU experiments using the distilled water, and FC-3284 spray cooling systems. . . . . 96

A-1 Thermal resistance diagram for heat transfer model. . . . . 108

## CHAPTER 1

### Introduction

This introductory chapter describes the problem that this research project aims to solve and frames that problem in the current context of the field. First, the objectives of this research are developed by analysing the problems and challenges involved with processor cooling. This development culminates with a refinement of the main objective into specific objectives. Subsequently, the academic literature, both directly and indirectly, concerned with processor cooling is reviewed. Finally, a concise overview of the remainder of this thesis is given.

#### 1.1 Background, Motivation and Overall Objectives

Computer systems are an essential and ubiquitous part of modern society. The enormous amount of electrical power these systems consume, which is ultimately converted to internal energy, has created an entire field dedicated to electronics and computer cooling. To put in perspective the sheer scale of power consumed and, therefore, the importance of electronics cooling, consider that, in 2014, data center electricity consumption alone constituted 1.8% of total U.S. electricity consumption (Shehabi *et al.* 2016). Furthermore, it is estimated that 22%–54% of the input power to a data center is required just to power the cooling systems (Tschudi *et al.* 2003).

The electronic components of a computer system convert inputted electrical energy into internal energy during normal operation. This energy must be removed efficiently; otherwise, the temperature of these components may rise to the point that they overheat and no longer function properly. Processors are the most vital components of a computer system since

they perform the actual computations. As would be expected, the processors also generate<sup>1</sup> the most heat<sup>2</sup> .

Most of the heat generated by a processor originates from the switching of transistors. Modern processors contain more than a billion transistors, and these transistors are concentrated in cores, which reside in the processor die. The heat emanates from the backside of the die, mainly from the small areas corresponding to the locations of the processor cores. Typically, the small backside area of the die is thermally connected to an integrated heat spreader (IHS) either by soldering or by application of a thin layer of thermal interface material (TIM). The main purpose of the IHS is to provide a larger area for the processor heat to spread laterally, while also protecting the die.

The larger area provided by the IHS is still inadequate, in terms of cooling area, for transferring the large amount of heat emanating from the processor to the environment. Therefore, some form of heat sink is required that (thermally) connects to the IHS. This thermal connection is created by a thin layer of TIM between the exposed, flat surface of the IHS and the bottom surface (cold plate) of the heat sink. The selection of TIM can have a meaningful impact on the overall thermal resistance of the cooling system; therefore, factors such as a TIM's intrinsic thermal conductivity, its thickness, and the thermal contact resistance to its two contact surfaces must all be considered during its selection (Mahajan *et al.* 2006). The remainder of the heat sink is typically designed to increase the overall heat transfer area and allow a cooling fluid to pass over it, thereby dissipating the heat

---

<sup>1</sup> Energy is neither created nor destroyed. However, it is common convention in this field to abbreviate the transformation of energy from one form to another to words synonymous with creation and destruction, such as generation and consumption, respectively. This convention is commonly used when referring to internal energy “generated” by the processor.

<sup>2</sup> Heat, in the strictest sense, refers to the transfer of energy from one body to another due to a temperature difference. However, it is common practice in this field to use the word “heat” to refer to internal energy that is either about to or is currently transferring away from a body. This convention is also adopted in this thesis.

by convection. In air-cooling heat sinks, the heat is removed by forcing air over numerous fins connected to the cold plate. In the more modern liquid cooling heat sinks, the heat is removed by pumping a cold liquid coolant through the heat sink. The internal energy imparted to the liquid coolant is eventually transferred to the surrounding air when it passes through a radiator.

Processors started out as fairly large components with relatively few transistors so low heat flux air cooling technologies were sufficient for a very long time. Gradually, processors shrunk and the number of transistors in them exponentially increased, but improvements in their thermal efficiency adequately limited any increases in heat density such that improved cooling technologies, such as liquid cooling, became necessary only in some cases. However, the constant miniaturization of electronic devices has brought upon a phenomenon known as self-heating, which is resistive heating due to phonon-electron scattering, and this heating cannot be limited by a more thermally efficient design (Garimella 2008). Therefore, further miniaturization and transistor count increases will only increase the heat density of processors and necessitate the advent of more advanced, higher heat flux processor cooling technologies.

While it is clear that more advanced processor cooling technologies will be eventually required, there nevertheless currently exists a significant need for improved cooling in high-performance computer systems. These high-performance systems are used in a variety of applications, including engineering design, gaming, and simulation. For example, in the finance industry, the success of a high-frequency trading scheme is dependent on the speed of the computer system. The ultimate speed of a computer system is chiefly dependent on the speed of the main processor, which is typically the central processing unit (CPU).

The most common indicator of a CPU's speed is its clock speed, or CPU frequency, which is a measure of the number of clock cycles per second. Although, it should be noted that this is only one of many factors that determine how fast a computer performs in a given application. A CPU is shipped with a default clock speed, which is considered the maximum recommended speed for normal operating conditions. With this limit on clock

speed, the CPU manufacturer can specify a thermal design power (TDP) for the CPU, which characterizes the maximum amount of heat that it generates while it is executing a typical set of realistic applications. The TDP is a useful metric when selecting a suitable cooling system for the CPU. However, there are methods of increasing the CPU clock speed, also known as overclocking, above its default clock speed. Overclocking is typically not recommended by CPU manufacturers because the CPU can be damaged or malfunction if the overclocking is performed without proper care. The increased rate of transistor switching also increases the rate that heat is generated, and if this additional heat is not transferred away efficiently from the CPU with an appropriate cooling system, the CPU overheats. Therefore, improvements to a processor cooling system have the added benefit of allowing an existing processor to operate at higher speeds.

Furthermore, processors are fast approaching their transistor density limit due to quantum mechanical limitations of silicon-based electronics, such as quantum tunnelling (Morton *et al.* 2011). Once this limit is reached, improvements in processor performance will only come by increasing the transistor-switching rate, which is limited by the effectiveness of the cooling system. Thus, the majority of processor performance improvement may soon arise from improvements to the processor cooling technologies rather than advances in processor technology. It is possible that this processor performance limit will eventually be overcome with development of quantum computing technologies; however, quantum computing is only superior to classical computing in certain computations and is not yet a readily available technology. Furthermore, like binary processors, quantum processors need to be cooled as well.

To ensure their continued performance, processors are cooled to maintain the core temperatures under a processor-dependent limit, which is usually between 75 °C and 150 °C. In fact, many processors monitor their thermal state and have a fail-safe mechanism that shuts down the processor once temperatures rise above that limit. This limit typically indicates

the point at which the processor may start to malfunction, or, in some cases, become damaged. Furthermore, lower processor temperatures are desirable because processors become exponentially less reliable with increased temperature. More specifically, it has been shown that a 10 °C increase in maximum processor temperature can decrease its reliability by as much as 50% (Peterson 1990).

Since most of a processor’s electronic activity is concentrated in the cores, these areas generate the most heat in the processor die. The cores of high performance processors of today can generate around 150 W/cm<sup>2</sup>, and the remaining background of the processor generates about 20 W/cm<sup>2</sup> (Sharma 2015). The higher heat flux in these areas has led them to be termed “hotspots”. These hotspots also tend to create significant temperature gradients across the surface of a processor die that can adversely affect operation. Thus, device reliability can be substantially improved using a cooling method that also maintains temperature uniformity across the die’s surface (Bostanci *et al.* 2009).

It is not surprising that other issues related to processor cooling exist. For example, one subtle problem is related to the processor surface not being completely flat and can even be warped due to differences in the coefficients of thermal expansion of the die and the package substrate (Prasher 2006). Verma *et al.* (1998) showed that the processor could become highly convex at an elevated temperature, which can reduce heat transfer due to an effectively smaller contact area. Moreover, Garimella (2008) identified multiple areas that he believes require further improvement to push the advancement of processor cooling. The areas requiring improvement include: low thermal resistance materials; interface interactions; electrothermal integration; and passive thermal management.

There is an extensive body of academic literature related to advanced cooling methods for dissipating heat from a flat surface. These cooling methods are thus also applicable to electronics and processor cooling. However, little research has been dedicated to the performance of these cooling technologies with actual processors. Furthermore, very few

commercial processor cooling products employ any of these advanced cooling methods. This indicates that this area of research still has a need for advancement.

The ultimate objective of this research is to improve the performance of processors by advancing the field of processor cooling. The first steps in achieving this objective and the major concerns of the present research are the design, construction and testing of an experimental apparatus for testing advanced processor cooling methods with both actual computer processors and flat surface heaters that simulate a processor. Initial testing with the highly-controlled, simulated processor establishes the heat transfer characteristics of the tested method, while the real processor testing determines the maximum processor performance attainable using that method. To validate this experimental apparatus and make an initial attempt at improving processor cooling, a potentially novel processor cooling device using an advanced cooling method is designed, built and tested using the experimental apparatus.

## **1.2 Specific Objectives**

This project focuses on the cooling of a CPU as it is the most common computer processor. The overall objectives of this work are subdivided into the following specific objectives:

- i) review advanced cooling methods and select the method with the best potential of achieving improved CPU cooling;
- ii) design and build an experimental apparatus, along with a heater that can simulate the CPU, for testing the designed CPU cooling device as well as future devices;
- iii) design and build a novel CPU cooling device based on the cooling method selected;
- iv) validate the system by conducting proof-of-concept testing of the cooling device with the simulated CPU; and
- v) validate the system by conducting proof-of-concept testing of the cooling device with the real CPU.

### 1.3 Literature Review

Advancing the field of processor cooling is a broad goal that first requires a similarly broad base of the current state of the field. This review of prior art begins with a brief overview of flat surface heat transfer, which is the basic form of heat transfer associated with processor cooling. Subsequently, an extensive examination of advanced cooling methods with applicability to processor cooling follows. Finally, various techniques and materials used to enhance the advanced cooling methods are investigated.

#### 1.3.1 Flat Surface Heat Transfer

The heat from a processor emanates from the back, nominally flat surface of the processor die, which represents a small area on the order of  $1 \text{ cm}^2$ . Furthermore, the majority of this heat is released from the processor cores, which represent only a small fraction of the area. In most cases, the processor die has an attached IHS that spreads this heat over a larger flat surface. Whether cooling the processor IHS or the processor die directly, the cooling method must be designed for or adapted to flat surface heat transfer.

Flat surface heat transfer is exceedingly common; therefore, many flat surface cooling methods exist, and the field is advanced. In most cases, the two most important parameters in surface heat transfer are the (surface) heat flux and the surface temperature. The heat flux quantifies the heat transfer rate per unit area. The surface, or wall, temperature generally refers to the temperature of the surface undergoing the heat transfer, averaged over its area.

When it comes to higher order heat fluxes, heat flux ranges are commonly defined by the following classification system (Ebadian and Lin 2011):

- High heat flux (HHF):  $10^2$ - $10^3 \text{ W/cm}^2$
- Ultra-high heat flux (UHF):  $10^3$ - $10^4 \text{ W/cm}^2$
- Extreme heat flux (EHF):  $>10^4 \text{ W/cm}^2$

Modern high-end processors produce heat in excess of  $100 \text{ W}$ , and since this heat emanates from surfaces with areas on the order  $1 \text{ cm}^2$ , the cooling methods reviewed should be capable of heat fluxes on the order  $100 \text{ W/cm}^2$ . Using the above classification system, a

heat flux of this magnitude is referred to as a high heat flux, and this is the heat flux range of primary interest for this review.

Heat flux is a function of surface temperature for any cooling method, and the surface temperature normally increases with increased heat flux if all other conditions remain constant. The heat transfer function of how the heat flux changes with surface temperature for a specific cooling method is as important as the maximum heat flux possible with that cooling method. However, the published research on higher heat flux cooling methods tends to suffer in terms of comparability, mainly due to the absence of the full heat transfer function. For example, many of the reported heat fluxes, especially maximum heat fluxes, are obtained without recording the surface temperature. Therefore, a certain scrutiny is required when analyzing the current state of the art and any reported performance values. This scrutiny is particularly important when evaluating the applicability of a method to processor cooling, which must maintain the maximum processor temperature below a certain limit. Nevertheless, studies with temperatures above those limits are still worthy of investigation because they give a better understanding of the full capabilities of a particular cooling method.

The heat transfer function also reveals the maximum heat flux a cooling method can achieve before failing. For liquid or two-phase cooling methods, the maximum heat flux is referred to as the critical heat flux (CHF). As the temperature of a heated surface is increased, the heat transfer coefficient eventually stops increasing and starts to greatly decrease. This is caused by the onset of dryout, and the heat flux at this point is known as the CHF. Dryout indicates that a vapour blanket has formed on the heated surface and is acting as an insulator. This greatly increases the surface temperature, which usually leads to failure (Betz *et al.* 2011).

### 1.3.2 Cooling Methods

The four main classifications of electronics cooling methods are air cooling, liquid cooling, two-phase cooling, and solid-state refrigeration. Although advancements in air cooling methodologies that increase their heat flux (e.g. branching radial fins, skived heat sinks,

ionic prime-movers, active cooling with piezo-actuated jets, and various hybrid systems, (Garimella 2008), air cooling approaches are not covered in this review since natural and forced air convection cooling is generally unable to achieve heat fluxes greater than approximately  $1 \text{ W/cm}^2$  (Ebadian and Lin 2011). This review therefore focuses on the remaining three categories; however, liquid cooling and two-phase methods are covered together since these approaches can typically be operated in both a single or two-phase system.

### **Liquid and Two-Phase**

Liquid cooling methods are currently the most popular advanced processor cooling methods. Two-phase methods are seldomly used in practice; however, two-phase cooling has the most potential for high heat transfer rates (Ebadian and Lin 2011). The methods that utilize two-phase cooling almost exclusively use liquid-vapour phase change due to the enormous amount of energy associated with the latent heat of vaporization. However, solid-liquid phase-change materials (PCMs) have been shown to be suitable in electronics applications in which the heat given off by the source must be stored in a compact form factor rather than transferred to the environment because they only undergo approximately a 10% or less change in volume (Sharma *et al.* 2009). The liquid and two-phase cooling methods reviewed herein are microchannels, jet impingement, sprays, heat pipes and vapour chamber, and liquid immersion.

**Microchannels.** Microchannels have been extensively researched in response to the ever-increasing heat fluxes generated by new electronic technologies. A microchannel cooling system consists of an array of channels, with hydraulic diameters in the micrometer range, that are side-by-side on the heated surface, and through which coolant is pumped as single or two-phase flow. Microchannels take up little space on the back of a processor and thus are particularly well-suited for multichip integration (Ebadian and Lin 2011). Moreover, they are able to remove large heat fluxes given that the associated convective heat transfer coefficients are very high due to the tiny hydraulic diameter of the channels. However, microchannels

do have shortcomings, including large pressure drops, temperature nonuniformities, coolant leaks, liquid maldistributions, and large temperature gradients along the direction of the flow (Ebadian and Lin 2011; Sung and Mudawar 2009). Furthermore, the large pressure drops can necessitate large pumping powers since the pumping power is directly proportional to the pressure drop across the channel.

Tuckerman and Pease (1981) conducted one of the first notable experiments using microchannels and proved that HHFs are possible in the compact spaces of microelectronics by achieving a maximum heat flux of  $790 \text{ W/cm}^2$  in their experiments. This heat flux was attained with a maximum temperature rise of  $71 \text{ }^\circ\text{C}$  in the substrate above the input water temperature. Tuckerman and Pease also noted that the convective heat transfer coefficient increases with decreasing hydraulic diameter for a fully developed laminar flow.

Hirshfeld (2006) obtained an even larger heat flux ( $1,500 \text{ W/cm}^2$ ) also using microchannels. However, this increase in heat transfer was coupled with a pressure drop of about 3.75 bar and a maximum temperature of  $450 \text{ }^\circ\text{C}$ . Moreover, Hirshfeld estimates that the maximum heat flux removal is likely higher than  $1,500 \text{ W/cm}^2$  since his experiment was limited by the loss in strength of his aluminum heat sink at those high temperatures.

Mudawar (1999) achieved an extreme CHF of  $27,000 \text{ W/cm}^2$  using subcooled water, a very high mass flux ( $120,000 \text{ kg/m}^2\text{-s}$ ) and small tube diameter (0.406 mm). However, this was only possible with an enormous pressure drop of 129 bar. Theoretically, higher mass fluxes and smaller tube diameters are not even attainable because the inlet pressure would approach the critical pressure. Mudawar also notes that increasing mass flux, increasing subcooling, decreasing tube diameter, and decreasing heated length-to-diameter ratio are methods of increasing ultra-high CHF. However, these approaches are coupled with drawbacks, such as increased pumping power.

Sharma (2015) developed a hotspot-targeting microchannel design that used a passive, energy saving approach by optimizing the microchannel geometry and flow rate distribution. He was able to achieve a maximum temperature nonuniformity of  $3.7 \text{ }^\circ\text{C}$  on the processor

surface with heat fluxes of  $150 \text{ W/cm}^2$  and  $20 \text{ W/cm}^2$  for the hotspots and background respectively. His design employed certain guidelines to focus on decreasing the nonuniformity in die temperature due to hotspots. These included:

- having the incoming cool liquid first impinge on hotspots and exit from the background areas;
- minimizing the microchannel wall width to reduce the pressure drop; and
- using the finest channels over the hotspots.

Two-phase flow is able to achieve much greater temperature uniformity, which greatly aids in the prevention of local hotspots. Generally, in the beginning of single-phase flow, the surface temperature is linearly dependent on the heat flux. Surface temperature drops rapidly at the onset of boiling, but it then starts to increase as more vapour is developed and the flow approaches CHF conditions. High temperature nonuniformity in the transverse direction can result from boiling in microchannels with axially nonuniform heating (Bogojevic 2011).

Lee (2009) utilized high inlet subcooling, high mass velocities, and small length-to-diameter ratios to ensure that bubbly flow endures much longer down the channel, and CHF, deemed Departure from Nucleate Boiling (DNB) CHF, begins when the bubbles form into a vapour blanket. In Lee's experiments, CHF's of around  $700 \text{ W/cm}^2$  were realized, and Lee predicts that with better testing equipment, CHF's above  $1,000 \text{ W/cm}^2$  could be attained.

A theoretical study of the effect of spatially nonuniform heat generation from a processor cooled by a two-phase microchannel heat sink has been conducted by Koo *et al.* (2002). Koo determined that minimization of pressure drop is the most important design consideration concerning microchannel heat sinks. The area of the processor generating more heat should be placed near the channel outlets because this minimizes pressure drop and surface temperature. Furthermore, the study found that the microchannel width should be as

large as possible; however, it was noted that wider channels might experience instability in performance due to the lateral variation of the flow.

In a theoretical design experiment of a refrigeration system with two-phase flow through microchannels, Marcinichen (2010) determined that for a coolant evaporating temperature of 60 °C, a subcooling temperature of 5 °C and an outlet vapour quality of 30%, the predicted CHF was 141.2 W/cm<sup>2</sup>, which is 2.2 times higher than the maximum heat flux using fins and air cooling. It was also found that increasing the vapour quality beyond 30% greatly decreased the CHF. The lower outlet vapour quality is also beneficial because the fluid can then be used to cool the remaining lower heat flux components of the server before entering the condenser.

Gong (2011) conducted experiments that explored the benefits of using wavy microchannels as opposed to straight channels. Most notably, the experiments found that an increase in thermal performance of up to 55% compared to a straight wall configuration was possible. The study also determined that the serpentine wavy configuration was most effective, and larger amplitude and smaller wavelength channels performed best.

Significant increases in bandwidth can be achieved with three-dimensional space (3D) integration of processors (Brunschwiler *et al.* 2008). In 3D processor stacks, the maximum heat flux interlayer cooling capability is strongly dependent on the heat transfer area. Brunschwiler found that the cooling capability dropped from values above 200 W/cm<sup>2</sup> for a 1 cm<sup>2</sup> area to less than 100 W/cm<sup>2</sup> for a 4 cm<sup>2</sup> area. Koo *et al.* (2005) conducted a similar study and predicted, with theoretical calculations, that his cooling strategy would be able to achieve heat removal up to 135 W/cm<sup>2</sup> within a 3D circuit architecture with a maximum temperature of 85 °C.

**Jet Impingement.** Jet impingement is another popular cooling technology when it comes to cooling large heat fluxes. It works by pumping a jet stream of coolant through a nozzle, which impacts the heated surface, typically normally. It is also common to have multiple jets spatially arranged in an array to cover the heated surface more uniformly. Massive

heat transfer coefficients are possible in the impingement zone. Also, jet impingement generates a relatively small pressure drop. However, large temperature gradients occur away from the impingement zone; therefore, multiple jets are typically required (Sung and Mudawar 2009).

Silverman (2006) was able to achieve ultra-high heat flux ( $2,000 \text{ W/cm}^2$ ) cooling, over an area of  $1 \text{ cm}^2$ , with a relatively low pressure differential of 1 bar. This was accomplished using GaIn liquid metal, with a melting point of  $15.7 \text{ }^\circ\text{C}$ , at a low jet impingement velocity of  $4 \text{ m/s}$ .

Fabbri *et al.* (2003) conducted experiments to compare spray against jet impingement cooling. For equal pumping power and a wall to liquid temperature difference of  $76 \text{ }^\circ\text{C}$ , the jets cooled heat fluxes of  $240 \text{ W/cm}^2$ , and the spray only achieved  $93 \text{ W/cm}^2$ . However, at matching flow rates, the spray was better than all jet configurations.

In the design of a cooling finger for a nuclear fusion cooling application, Koncar (2010) investigated the effect of nozzle size on jet impingement cooling using helium as the coolant. The study found that the maximum cooling finger temperature reduction ( $27 \text{ }^\circ\text{C}$ ) was obtained using equal sized nozzles of  $0.6 \text{ mm}$  diameter in a cartridge design. Decreasing the diameter of nozzles, specifically the center nozzle, caused the maximum cooling finger temperature to decrease linearly, while the pressure differential increased linearly.

A few studies have attempted to obtain maximum possible heat flux. For example,  $40,000 \text{ W/cm}^2$  fluxes have been successfully handled using subcooled water jet impingement on a miniscule area ( $\sim 1 \text{ mm}^2$ ) (Liu and Lienhard 1992). However, these methods are not applicable to most applications due to various constraints, such as size and material use.

A microchannel and jet impingement hybrid system was designed and tested by Sung and Mudawar (2009). The system attempted to eliminate some of the weaknesses exhibited by each individual method. The jets were used to inject the coolant, HFE 7100, into the microchannels at multiple locations along the channels. This design prevented large temperature gradients from forming in the flow direction and also lowered the pressure drop. Two

jet configurations were tested: circular and slot. It was determined that circular jets were able to handle higher heat fluxes due to the repeated pattern of bubble growth and collapse between jets. The study also determined the optimal microchannel height required for this unique solution. The height had to be small enough for the jet coolant to actually impinge on the surface but tall enough not to decrease the cooling surface area provided by the walls too much. The maximum heat flux cooled by the hybrid system was  $1,127 \text{ W/cm}^2$ , but the equipment in use, rather than CHF, limited this value.

**Sprays.** Spray cooling technology is similar to jet impingement in that it injects a coolant through, typically, an array of nozzles at a heated surface. In contrast to jet impingement, spraying usually attempts to maximize heat dissipation by utilizing the liquid-vapour phase change process by atomizing the coolant before it contacts the heated surface. Multiple mechanisms of cooling are at work in this method. Boiling through surface and secondary nucleation, evaporation at the liquid film surface, and forced convection from the actual impingement of the coolant on the surface all play a role in absorbing the heat load (Bostanci *et al.* 2009). Spray cooling has also been deemed very attractive due to its ability to deal with large heat fluxes, while maintaining temperature uniformity across the heated surface (Bostanci *et al.* 2009).

Cooling of a heated surface by spraying technology has been under investigation for decades. Pais (1992) was able to cool a large heat flux of  $1,200 \text{ W/cm}^2$  in the 1990's using spray cooling. This was in part accomplished by modifying the spraying surface profile by polishing it with emery paper of grit size  $0.3 \mu\text{m}$ .

Bostanci *et al.* (2009) used an antifreeze, with a  $90 \text{ }^\circ\text{C}$  boiling point, as a coolant and removed a  $400 \text{ W/cm}^2$  heat flux from a heated surface, while maintaining the surface temperature below  $105 \text{ }^\circ\text{C}$ . The relatively small amount of surface superheating indicates a very small thermal resistance from the spraying method.

The effects of surface enhancement and spray inclination angle were analyzed by Silk (2006). Cubic pin fins, pyramids and straight fins were tested as surface enhancements

with a flat surface used as a baseline, and all surface enhancements had an improved CHF compared to the flat surface. The spray inclination angle was tested for angles between  $0^\circ$  and  $45^\circ$ , and in all cases an inclination improved the CHF over the  $0^\circ$  baseline. The best configuration was straight fins at an inclination angle of  $30^\circ$ , which resulted in a maximum heat flux enhancement of 75% over the base case.

**Heat Pipes and Vapour Chambers.** A heat pipe is a passive heat transfer device that moves internal energy by means of two-phase coolant fluid flow. Heat pipes come in many shapes and sizes, but they can all generally be broken down into three sections. The evaporator section is where the heat is absorbed and the coolant vaporized. Consequently, there is also a condenser section in which the coolant condenses and releases the absorbed heat. The last section, located in between the previous sections, is called the adiabatic section, and it is simply there to create a connection for fluid flow between the evaporator and condenser sections. Different types of heat pipes have different methods of getting the liquid from the condenser back to the evaporator. Many methods make use of a wick system that is dependent on capillary action. This simple device can be designed in a way such that it can displace heat from the heat source to anywhere in an efficient manner. A heat pipe can actually transfer heat from one place to another hundreds of times more efficiently than a solid copper device of the same form (Mochizuki *et al.* 2009).

Heat pipes have been used with computers and processors for over two decades. In portable notebook computers, miniature heat pipes are used to displace processor heat to a more convenient heat dissipation area. Heat pipes can also be used in a remote heat exchanger configuration in which the condenser is placed away from the heat source and is given multiple fins for dissipating heat to the environment with the aid of a fan. In desktops and servers, a heat transfer device developed with the same principles of a heat pipe, called a vapour chamber, is more commonly used (Nguyen *et al.* 2000).

Vapour chambers are similar to heat pipes in that they too are heat transfer devices that move internal energy by means of two-phase coolant fluid flow. However, vapour chambers

are designed to improve heat transfer in a second dimension as well. The two-dimensional heat flow has the advantage of efficiently spreading heat laterally, which improves temperature uniformity.

Vapour chambers are therefore perfect for spreading the heat of a processor die's non-uniform heat flux distribution. In fact, the use of a vapour chamber variant, known as a microchannel vapour chamber IHS (MVC-IHS), instead of a typical solid copper IHS indicated a performance increase of 40% (Mochizuki *et al.* 2009). If fins are added to the vapour chamber, it becomes a full-fledged heat sink.

**Liquid Immersion.** Direct liquid immersion is appealing because the entire heat-generating package can be fully submerged in a high sensible heat liquid, such as water or oil. The obvious challenge in this type of design is ensuring that there is no possibility of damage to the electronic components of the package. Haywood (2015) determined that direct CPU submersion in oil lowered the drop in temperature between the heat source and the discharge line by 81% compared to a traditional water-cooling system in the same conditions. This indicates a much lower thermal resistance in the oil system and a significant boost in efficiency. It should be noted that dielectric oils with very low electrical conductivity were used to avoid any disruption in the computer operation. However, Haywood's experimental system had multiple problems, such as pump failure and air entering the tubing lines, which caused premature disruptions in the experiments and prevented the longer running times that are necessary to achieve accurate and reliable data in this type of experiment. Although liquid immersion cooling is attractive in certain ways, it is typically incapable of handling high heat fluxes due to the lack of significant convection.

### **Solid-State Cooling**

Interest in solid-state cooling for electronics applications has spiked in recent years due to its ease of control, and lack of moving parts and fluids. Much of the current research is

focused on improving efficiency and decreasing cost. Thermoelectric cooling and solid-state cooling with caloric materials are examined below.

**Thermoelectric Materials.** Thermoelectric cooling uses a technique, known as the Peltier effect, in which an electric potential is created between two surfaces in order to induce heat flow. Thermoelectric cooling is attractive in many applications because it is instantly variable solid-state refrigeration. Sub-ambient temperature cooling is possible without the need for the typical, complex refrigeration cycle. Two more advantages of thermoelectric cooling are its ability to target specific locations, and to do so on command. Therefore, it is the perfect option for hotspot mitigation, and it can offer massive improvements in efficiency by improving the uniformity of the heat transfer profile of dies and by being able to turn off when not needed (Chowdhury 2009).

ZT is the thermoelectric figure of merit, which is a quantity used to characterize a thermoelectric material's performance relative to other materials, and it takes into account thermal conductivity, electrical resistivity and the Seebeck coefficient. To validate the use of thermoelectric cooling, the current material ZT values must be significantly improved, most likely by reducing the lattice thermal conductivity by scattering. Furthermore, the heat flux pumping capacity of established coolers must also be increased (Chowdhury 2009; Garimella 2008). In one study, Chowdhury (2009) used  $\text{Bi}_2\text{Te}_3$  superlattice-based thin-film thermoelectric coolers to cool a silicon processor. The study showed local spot cooling of  $15\text{ }^\circ\text{C}$  at a heat flux of  $1,300\text{ W/cm}^2$ , proving that thermoelectric cooling may be feasible with processors.

**Caloric Materials.** Another solid-state cooling method utilizes caloric materials. Like vapour-compression refrigeration, caloric materials also use a cycle in their cooling process. This is the main differentiator between caloric materials and thermoelectric materials, which do not require a cycle to operate. The cycle is called the ferroic cooling cycle. It begins by inducing a form of transition in the material, which raises the temperature. The material subsequently releases heat due to the temperature difference with the environment. Next,

the force used to induce transition in the material is released, which causes the material to cool to a temperature below the initial temperature. This allows the material to absorb heat from the load to return to its original state. Barocalorics are one type of caloric material that undergoes a phase transition due to a change in hydrostatic pressure. This means that common liquid-gas refrigerants are actually considered barocaloric materials (Takeuchi and Sandeman 2015). Magnetocaloric, electrocaloric and elastocaloric materials all undergo similar transitions to produce a cooling effect.

Magnetocaloric cooling uses a magnetic field to transition a suitable material from dis-oriented magnetic domains to oriented domains, which raises the material temperature. Removing the magnetic field significantly lowers the temperature of the material, which allows it to then be used to cool. Electrocaloric cooling functions similarly to magnetocaloric cooling but instead uses an electric field, which causes a transition in electric polarization of the material. Elastocaloric cooling is different because it uses a mechanical stress rather than a field to induce transition. The change is in the crystalline structure of the material.

Like thermoelectric cooling, caloric cooling is specifically useful due to the fact that it makes sub-ambient temperature cooling possible without a cold reservoir. Caloric cooling also has the major advantage of having extremely high coefficients of performance (COP) in the range of 10 to 25 before taking into account the losses associated with putting them into a cyclic system. Furthermore, extremely high latent heat energy densities, on the scale of MJ/m<sup>3</sup>, are possible with these materials (Takeuchi and Sandeman 2015).

Caloric materials do have some immediate disadvantages as well. Although caloric materials have been studied for many years, the technology still requires further development to become more economically attractive. Also, systems that can use their cooling ability are more complex because they must operate cyclically. Furthermore, a cyclic system operating on a solid material inherently introduces fatigue issues. Lastly, the caloric cooling cycle for these materials typically operates in small temperature ranges, which may limit its usefulness for many applications (Takeuchi and Sandeman 2015).

### 1.3.3 Cooling Enhancement

While some of the enhancements for each individual cooling method have already been discussed, many more techniques, engineered coolants, materials and structures exist that are capable of enhancing the heat transfer of the various cooling methods. Nanofluids, nanostructures and surface wettability are a few of the more promising enhancements, and they are discussed below.

#### Nanofluids

Nanofluids consist of a base fluid mixed with nanoparticles that are between 1 nm and 100 nm in size. Common base fluids include water, ethylene glycol and oil, and common nanoparticles include copper and carbon nanotubes (CNT). Nanofluids were developed to increase thermal conductivity of typical heat transfer fluids, such as water. Thermal conductivity of nanofluids is influenced by factors such as temperature, particle size, dispersion and stability (Murshed 2009). The high specific surface area of the particles increases the heat transfer between the particles and the fluid. Nanofluids are also appealing for having adjustable properties, such as surface wettability (Saidur 2011). A nanofluid is most useful in its applications when the nanofluid exhibits uniform dispersion and stable suspension of the particles in the base fluid (Lee 2008).

Nguyen (2007) showed that a commercially available  $\text{Al}_2\text{O}_3$ -water nanofluid has a convective heat transfer coefficient that is 40% higher than distilled water in a microprocessor cooling application. It was also shown that the smallest sized nanoparticles (36 nm in diameter) and the highest tested particle volume concentration (6.8%) give the best heat transfer characteristics. The benefit to heat transfer of using smaller particles can be simply attributed to the fact that, at the same volume concentration, the smaller particles have a higher contact area for heat exchange (Nguyen 2007). The benefit of higher volume concentration can be similarly explained; however, Nguyen's experiment did not conduct any tests with concentrations higher than 6.8 vol% so it is unclear if increasing the concentration continues to yield better heat transfer coefficients. It is likely that increased concentration

eventually leads to problems in the flow of the liquid and pumping issues. In the low volume concentration range of 0.01%–0.3% for  $\text{Al}_2\text{O}_3$ -water nanofluids, the viscosity decreases significantly with increasing temperature, and the experimentally measured relation between concentration and viscosity is non-linear, which contradicts the accepted linear relation predicted by the Einstein viscosity model (Lee 2008).

Hassan (2015) showed that the use of nanofluids alternatively to water in a heat pipe results in improved heat pipe performance. The temperature of the evaporator was shown to be lower with nanofluids indicating better overall performance, and the temperature difference between the evaporator and condenser was shown to be up to 50% smaller compared to water, which points to a lower thermal resistance when using nanofluids. Hassan tested 1 vol% and 3 vol% nanofluids and found that the 3 vol% nanofluid performed best in every test. However, they also alluded to a problem with repeatability of the nanofluids experiments, which they believe is due to factors such as the poor stability of the nanofluids.

Saidur (2011) has identified some of the major issues that nanofluids are currently facing. Firstly, long-term stability of the nanoparticles' dispersion in the base fluid must be achieved before nanofluids can be reliably used. Furthermore, their mechanism of heat transfer is still poorly understood, which is also limiting researchers' ability to overcome the problem of nanofluids having a lower specific heat capacity compared to their base fluid. Finally, difficult production and high costs are limiting more widespread use of the engineered fluids.

## **Nanostructures**

Nanostructures have not received as much attention as nanofluids. However, nanostructures have incredible potential for use in TIMs or as heat spreaders due to their very high thermal conductivities. Hu *et al.* (2005) have shown that CNT-based TIMs have thermal conductivities up to  $83 \text{ W/m}\cdot\text{C}$ , which is one order of magnitude greater than commercial greases and phase change materials.

One of the current major issues facing nanostructures, such as CNT, is the need to reduce their contact resistance (Garimella 2008). Moreover, nanostructures are typically difficult to

integrate into a system effectively. Chemical vapour deposition (CVD) has been shown to be an effective method of applying CNT, with a high growth rate of 100  $\mu\text{m}/\text{min}$ ; however, CNT are not best used as heat spreaders (Garimella 2008).

Srikar *et al.* (2009) determined that the bouncing back of coolant droplets after impact on a heated surface, which is common in methods such as jet impingement or spray cooling, reduce heat transfer. Their solution was to apply a woven nanofiber mat to the heated surface. Instead of bouncing back, droplets spread out through the mat, remained close to the surface, and were able to evaporate. In one test, the impact of a droplet on a non-coated surface dropped the temperature from 60  $^{\circ}\text{C}$  to 41  $^{\circ}\text{C}$ , while the same impact on a nanofiber mat-coated surface dropped the temperature from 60  $^{\circ}\text{C}$  to 33  $^{\circ}\text{C}$ .

### **Surface Wettability**

The features of a heated surface can have various and important effects on the heat transfer characteristics of any cooling method. One surface feature that is alterable is the wettability of a surface, and it can be improved by applying a microstructure or nanostructure coating, or by changing the surface roughness (Betz *et al.* 2011). The wettability of a surface is described by the contact angle it has with the liquid. A hydrophilic surface has a contact angle close to zero, which corresponds to the liquid wetting the surface in a blanket fashion, while a hydrophobic surface has a contact angle much greater than zero, which causes the liquid to form into near-spherical droplets.

Liu (2006) conducted an experiment that determined the effect that a hydrophilic surface has on heat transfer for water jet impingement cooling. A  $\text{TiO}_2$  coating was applied to a copper surface and irradiated with ultraviolet light to make the surface superhydrophilic. The contact angle of the water was essentially reduced to zero, and this caused a large delay in the boiling incipience. This results in a CHF that is 50% higher than that of a common copper surface for both the subcooled water and the saturated water jet impingement variations. It should also be noted that superhydrophilic surfaces do not alter the effects on CHF of varying the impact velocity, the nozzle diameter and the subcooling of the coolant.

Experiments have also been conducted that investigate the effects of varying the wettability across the heated surface. Betz *et al.* (2011) devised a novel technique of putting a pattern of hydrophobic islands on a superhydrophilic surface, known as hydrophilic networks. This technique has been shown to be able to increase the CHF by 80% and double the heat transfer coefficient compared to a very hydrophilic surface.

#### **1.4 Overview of Thesis**

The remainder of the thesis is divided into an additional five chapters. Chapter 2 discusses the development of the experimental apparatus, including its fluid loop, two testing platforms, and the control and data acquisition system. Then, the design and prototyping of a novel processor cooling device is covered in Chapter 3. Subsequently, Chapter 4 outlines the methods followed to conduct the experiments, and Chapter 5 analyzes the results from those experiments. The thesis concludes in Chapter 6 with a review of the present work and plans for future work.

## **CHAPTER 2**

### **Experimental Apparatus**

The following chapter describes the experimental apparatus created for testing novel processor cooling devices, including the prototype described in the subsequent chapter. After an extensive review of the potential processor cooling methods, it was determined that an apparatus capable of testing liquid and two-phase cooling methods would have the greatest utility in terms of testing as many different cooling devices as possible. Therefore, the experimental apparatus needs a fluid loop, which is discussed in the first section of this chapter. The ensuing section is about the creation of the simulated CPU (heater) that is used during initial tests of a cooling device. Then, the real CPU and computer system used for final testing is described. In the final section of this chapter, the creation of the control and data acquisition system is presented.

#### **2.1 Fluid Loop: Design and Construction**

This section is concerned with the development of the fluid loop for the experimental apparatus. This fluid loop, in general, is similar to any closed liquid cooling loop used in current processor liquid cooling systems in terms of its major components. To maximize the types of cooling technologies that the apparatus can test, the fluid loop was designed to be compatible with both single-phase (liquid) and two-phase cooling devices. Typically, fluid loops capable of handling two-phase flow have added components, such as a condenser, in addition to the components of a liquid fluid loop. However, it has been determined that these additional components are not strictly necessary (e.g. the radiator can additionally perform the function of a condenser); thus, the same loop is used for both types of fluid flow.

In the first part of this section, an overview of the design is given alongside a diagram of the fluid loop. Next, the pump, radiator, fans, reservoirs, valves, sensors, tubing, and

fittings are described. Finally, the end result of the construction and assembly of the loop are presented.

### **2.1.1 Overview**

In a simple fluid loop for a processor liquid cooling system, a pump forces coolant out of the reservoir, into the CPU cooling device (typically a heat sink), then into the radiator, and finally back to the reservoir. The designed fluid loop (see diagram of figure 2-1) for the present experimental apparatus is similar in nominal operation but has some added complexities. The sensors are the first added complexities, and they include five inline temperature sensors, along with three pressure sensors. Furthermore, in addition to the main reservoir, a backup reservoir is included to intake coolant in the event of a problem in the main loop. Moreover, five valves are used to further control the flow of coolant and perform safety functions (described below). Lastly, two particle filters were added inline to trap any debris in the coolant before it enters into the CPU cooling device.

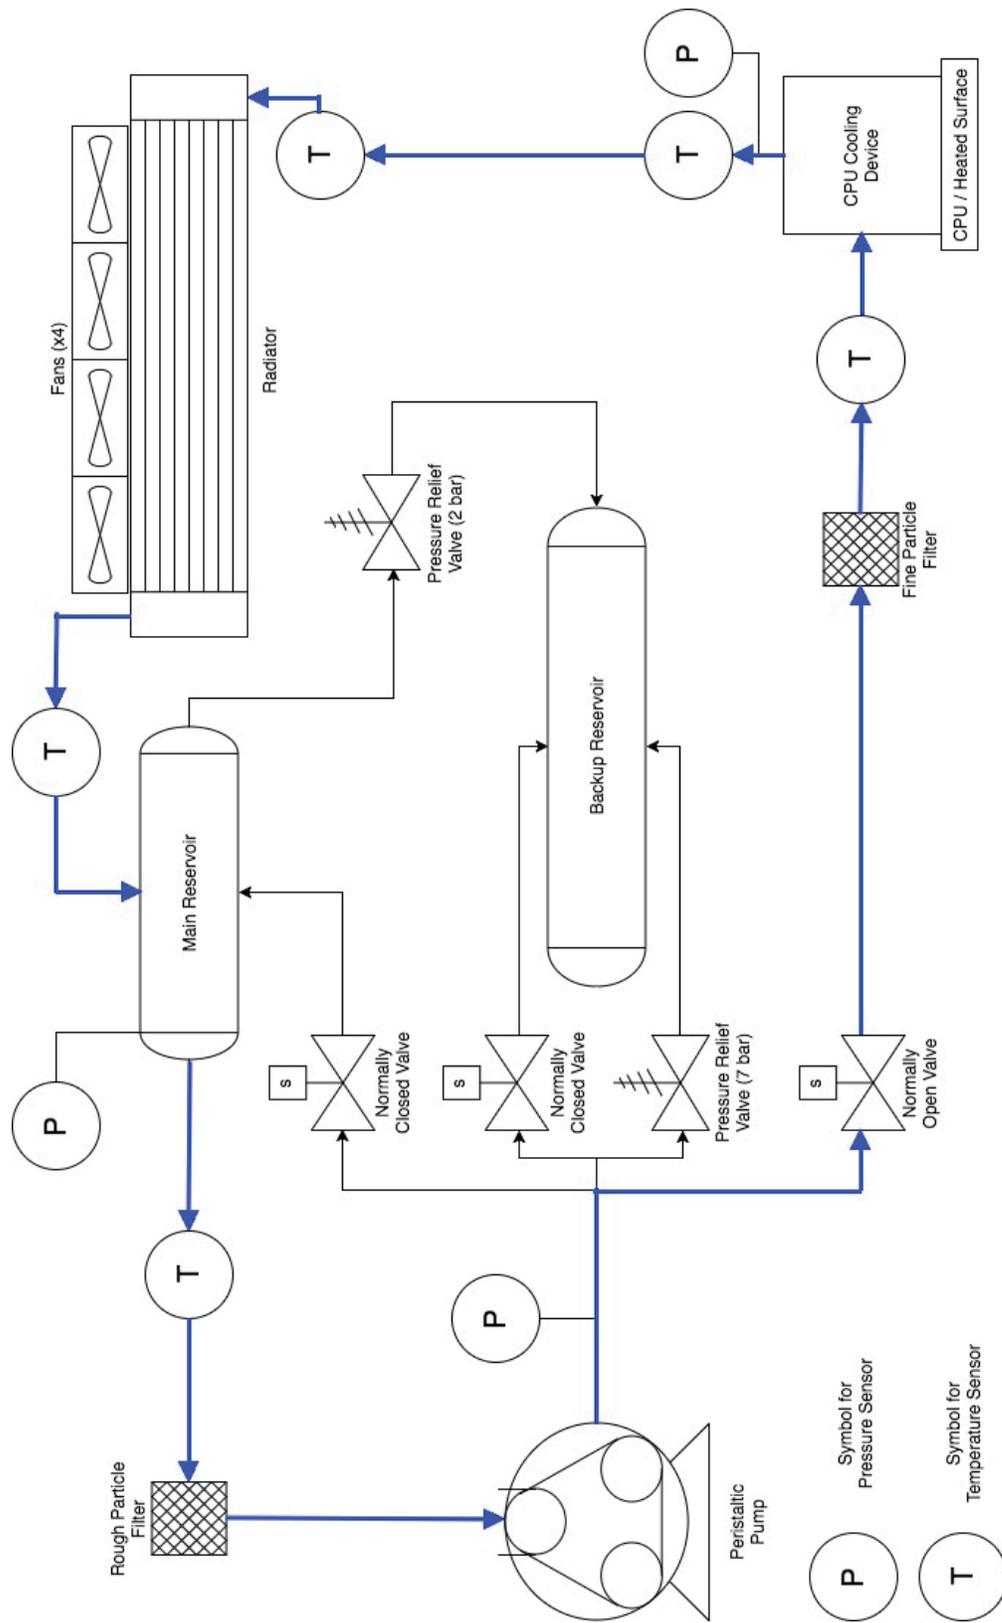


Figure 2-1: Experimental apparatus fluid loop diagram.

Liquid coolant flows out of the main reservoir, where its pressure is measured, and through a first inline temperature sensor and a rough particle filter before entering the pump. The pump pressurizes the coolant and forces it into a major junction, where its pressure is again measured. At this junction, the coolant has four different pathway options. The first option is towards a normally closed valve that leads back to the main reservoir. The second option is towards another normally closed valve that instead leads to the backup reservoir. The third option is towards a mechanical pressure relief valve that releases the coolant into the backup reservoir should the fluid pressure reach the maximum rated pressure of the pump (7 bar). The fourth and final option is through a normally open valve towards the CPU cooling device. This is the typical pathway that the coolant follows, and it passes through a fine particle filter and a second temperature sensor before reaching the cooling device. In the cooling device, the liquid coolant absorbs heat from the simulated or real CPU and potentially partially vaporizes. Then, the coolant exits the device through its outlet, where both a pressure sensor and a temperature sensor are located. The coolant then flows along a longer piece of tubing and through a fourth temperature sensor before entering the radiator. Forced convection of air through the fins of the radiator transfers heat away from the coolant and to the environment. Any vaporized coolant condenses, and the liquid coolant returns to near ambient temperature. It then flows through a final temperature sensor and back into the main reservoir. Should the pressure in the main reservoir reach unsafe levels (2 bar), a pressure relief valve releases the coolant into the backup reservoir.

### **2.1.2 Pump**

The pump is the most important component of the fluid loop because it controls the coolant flow rate, which dictates the flow through the CPU cooling device; therefore, it is essential that a pump with accurate flow rate control is used. A positive displacement pump was selected because they more accurately control the flow rate than do centrifugal pumps. A rotary-type positive displacement pump (specifically, a peristaltic pump) was chosen over

a reciprocating-type pump for its smoother flow. Furthermore, rotary pumps are better capable of handling vapour in the pump line, and thus require less supporting equipment.



Figure 2-2: Qdos 30 Universal+ Peristaltic Metering Pump from Qdos Brochure (2016).

The Qdos 60 Universal+ Peristaltic Metering Pump, from Watson-Marlow, was selected. Figure 2-2 shows the Qdos 30 version of the pump, which is identical to the Qdos 60 in appearance. It can output flow rates ranging from 0.1 mL/min to 1,000.0 mL/min, and it can do so with a resolution of 0.1 mL/min. Accurate flow rates in the lower range between 0.1 mL/min and 200.0 mL/min are not possible with a centrifugal pump; therefore, this peristaltic pump is ideal for CPU cooling devices that require low, accurate flow rates. Furthermore, the pump is capable of operating at its entire range of flow rates with discharge pressures up to 7 bar. Discharge pressures above 7 bar, either due to the characteristics of the cooling device or due to a blockage, damage the pump, so, as mentioned in the overview, a pressure relief valve rated for this pressure was added to the outlet junction of the pump.

The pump also has a variety of input and output connections. In addition to being able to control the pump manually with the display and buttons, the pump can also be controlled

with electrical connections to its input pins. The pump's run/stop function and flow rate are currently controlled by the control and data acquisition system of the experimental apparatus. Specific coolant doses and coolant recovery can also be controlled with its input pins. Furthermore, by connecting to the pump's output pins, the control and data acquisition system monitors its output parameters, including the run status, flow rate and alarm status.

### **2.1.3 Radiator and Fans**

Once the pumped coolant exits the CPU cooling device, the radiator and fans are responsible for transferring its absorbed heat to the environment. This is accomplished by forced convection of air through the radiator's fins using the fans. The goal of this assembly is to condense any coolant vapour and bring the liquid coolant back to ambient temperature so that it may return to the reservoir and once again be pumped through the system. The radiator and fans selected are standard components used in CPU liquid cooling systems; however, the oversized versions were chosen to ensure that they never limit the overall heat transfer capability of the experimental apparatus. XSPC's RX480 Quad 120 mm Radiator V3 was selected for this purpose. It has two rows of copper tubes (21 mm  $\times$  2 mm) and split fins spaced at 5 fins per cm. Four Orion Axial Fans (Part Number: OD1238-12HBXJ10A) from Knight Electronics were connected to the radiator. These fans are each capable of a maximum flow rate of 7.08 m<sup>3</sup>/min (or maximum static pressure of 368 Pa) at 6,000 RPM using 44.4 W of power. They require a 12 VDC power supply, allow pulse width modulation (PWM) control, and provide a tachometer connection.

### **2.1.4 Reservoirs**

The coolant exiting the radiator enters into a 440 mL EK-RES X3 250 reservoir (the main reservoir) made by ekwb. This reservoir is cylindrical in shape and has input and output ports at both ends of the cylinder. It's length is larger than its radius, which allows it to stand upright so that the pump can suck in liquid coolant from the bottom of the cylinder; thus, no vapour near the top can possibly enter into the pump line. Another advantage of

this specific reservoir is that the top and bottom ends each have numerous ports allowing multiple fluid loop and sensor connections.

The backup reservoir (EK-RES X3 400) is from the same family of ekwb reservoirs, but it has a greater volume (820 mL). This reservoir makes three connections to the fluid loop and its function is to take in coolant in the case of an emergency, such as an overpressured pump, so that the coolant is not released to the environment. Unlike the main reservoir, the backup reservoir lays horizontally so that coolant can freely flow down towards it.

### **2.1.5 Valves**

Valves control the direction of coolant flow and provide safety measures in the fluid loop. The first is a manual shutoff valve that is connected to the bottom of the main reservoir. It is normally closed, which allows drainage of coolant from the fluid loop. The remaining valves are pressure relief valves and solenoid valves.

One of the two pressure relief valves is rated for 7 bar and is placed between the pump outlet and the backup reservoir to protect the pump. The other is rated for 2 bar, is placed between the main reservoir and backup reservoir, and protects the radiator and main reservoir.

The solenoid valves are controlled electronically by the control and data acquisition system of the experimental apparatus. All three are connected to the major junction at the pump outlet. The first two are normally closed and are opened to send coolant back to either the main reservoir or the backup reservoir. The last valve is normally open and leads to the CPU cooling device. It is closed when flow to the cooling device must be stopped.

### **2.1.6 Sensors**

The fluid loop is equipped with multiple sensors that feed data to the control and data acquisition system. These include temperature, pressure and liquid level sensors.

## Temperature

The first temperature sensor is an ambient temperature sensor (DHT22-AM2302) that also measures relative humidity. This digital sensor is pre-calibrated and has an accuracy of  $\pm 0.5$  °C. It is useful for comparing the ambient temperature to the coolant temperatures in the fluid loop.

The five remaining temperature sensors are inline with the fluid loop and measure the temperature of the coolant as it passes through them. The first three in the fluid loop (outlet of main reservoir, and inlet and outlet of CPU cooling device) are XSPC G1/4 Sensors, and the final two (inlet and outlet of radiator) are Phobya G1/4 Sensors. Both brands are thermistors with negative temperature coefficients (NTCs). These sensors were not previously calibrated, so they were calibrated with a three-point calibration at 0 °C, 20 °C and 100 °C using the Steinhart-Hart model equation (Steinhart and Hart 1968):

$$\frac{1}{T} = A + B \ln(R) + C[\ln(R)]^3, \quad (2.1)$$

where  $T$  is the temperature in K,  $R$  is the resistance in  $\Omega$  across the thermistor, and  $A$ ,  $B$  and  $C$  are the Steinhart-Hart coefficients. This calibration found that the coefficients for the XSPC sensors are:  $A = 1.21 \times 10^{-3}$ ,  $B = 1.89 \times 10^{-4}$ , and  $C = 5.22 \times 10^{-7}$ ; and the coefficients for the Phobya sensors are:  $A = 1.57 \times 10^{-3}$ ,  $B = 1.21 \times 10^{-4}$ , and  $C = 8.40 \times 10^{-7}$ .

## Pressure

Three pressure sensors (BQLZR G1/4 Pressure Transmitter) were added to the fluid loop: one at the main reservoir, one at the pump outlet, and one at the CPU cooling device outlet. These pressure sensors are pre-calibrated with an accuracy of  $\pm 0.5\%$  of FS and a range of 0 to 10 bar.

## Liquid Level

The liquid level sensor (Koolance 70mm G1/4 Liquid Level Sensor) is installed in the main reservoir. It notifies the control and data acquisition system if the coolant level falls below

70 mm, which is approximately equivalent to 120 mL of coolant remaining in the reservoir. This indicates that there is either a leak in the system or too much coolant vapour.

### 2.1.7 Tubing and Fittings

The fluid loop and all of its components are connected together with tubing and numerous fittings. Tygon B-44-4X tubing was selected for its good chemical compatibility, ability to inhibit bacterial growth, and flexibility. The selected inner diameter for all tubing is 3/8" (9.525 mm) because it is a standard diameter. The tubing on the low-pressure side of the fluid loop (outlet of cooling device to inlet of pump) is clear, has an outer diameter of 9/16" (14.2875 mm), and is capable of withstanding pressures up to 2.1 bar. However, the tubing on the high-pressure side (outlet of pump to inlet of cooling device) is a high-pressure variation of the Tygon tubing that has an inner braid for reinforcement and outer diameter of 5/8" (15.875 mm). It is also capable of withstanding pressures of up to 11.7 bar.

The variety of fittings used to make the connections in the fluid loop came from many different vendors (Bitspower, XSPC, Alphacool, McMaster-Carr, etc.). The most common type of fitting used was the barb. Barbs, with the aid of a worm clamp, are used to convert the ends of a piece of tubing into threaded ends that can connect to other components. Another frequently used fitting is the quick-disconnect fitting, which is used on both ends of the CPU cooling device to allow easy and safe attachment to and removal from the fluid loop. Furthermore, many rotary barbs and rotary connectors were used to allow unstrained lateral movement of components and facilitate any modification to the loop by making it possible to remove a tubing line without having to cut it. The remaining fitting types include: 3-way connectors, 4-way connectors, many adapters, right-angle connectors, flow stoppers, etc. The fittings used were mostly of the G1/4 thread type since this is the standard thread used with CPU liquid cooling components. However, NPT thread type fittings were also required for certain components, such as some of the valves.

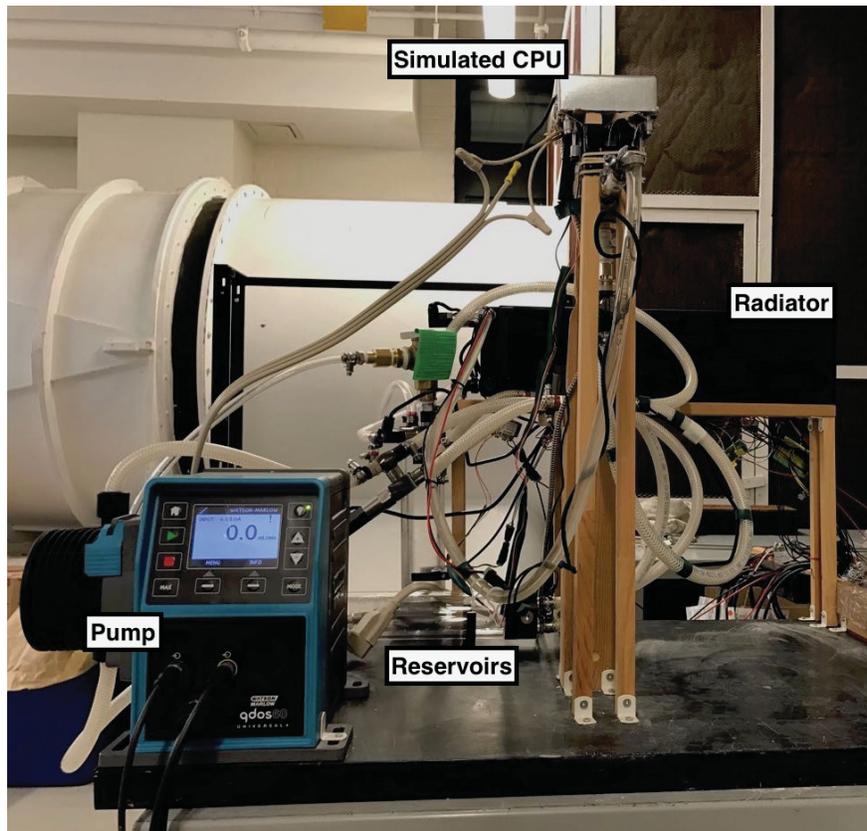


Figure 2-3: Experimental apparatus fluid loop - front view.

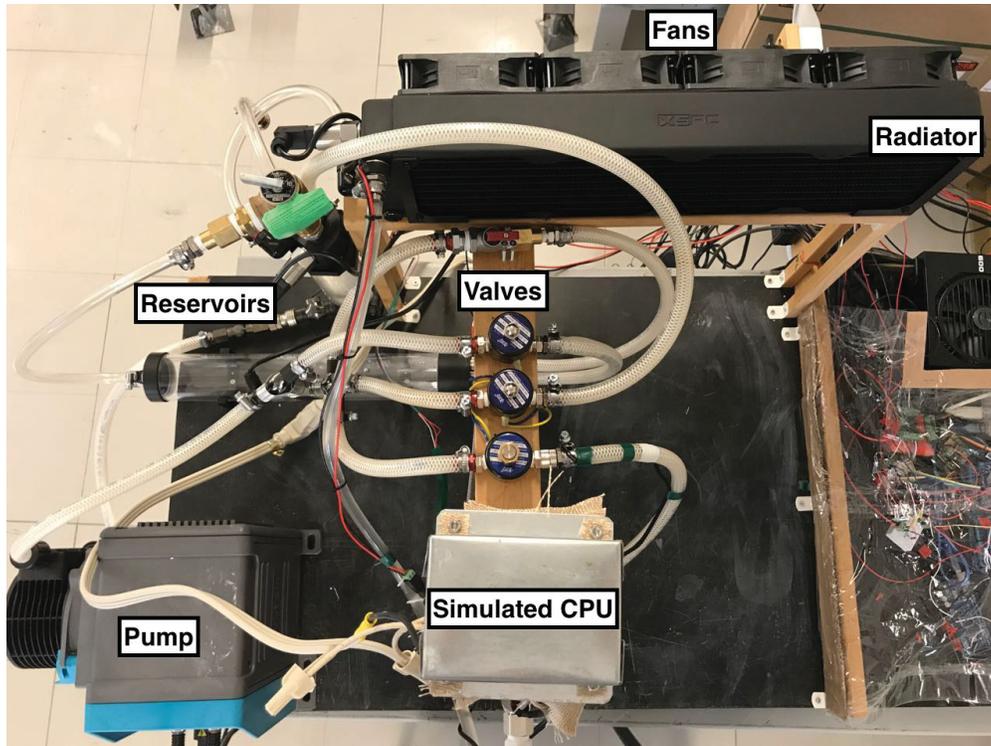


Figure 2-4: Experimental apparatus fluid loop - top view.

### 2.1.8 Construction and Assembly

The finished fluid loop is shown in figures 2-3 and 2-4. The entire system was built into a single wooden platform to ensure rigidity and stability, as well as to increase portability. The pump and backup reservoir are secured directly to the platform, which also has a series of stands built into it. The tallest stand seen in the front of figure 2-3 supports the simulated CPU and the cooling device so that liquid coolant flowing out of the device is always capable of flowing down towards the radiator regardless of the cooling method being tested. The radiator is built into the stand in the back at a height slightly higher than the inlet at the top of the main reservoir so that liquid coolant can flow towards the reservoir with the assistance of gravity as well. Referring to figure 2-4, a small platform has been built in between the front and back stand to further stabilize the stands and to secure the four valves connected to the junction at the pump's outlet. Lastly, the main platform space to the right of the front stand is reserved for the server system during the real CPU tests.

## 2.2 Simulated CPU: Design and Construction

When a CPU cooling device is connected to the fluid loop of the experimental apparatus, it is first attached to and tested using the simulated CPU. The simulated CPU is a resistance heater that mimics the overall heat generation and heat transfer of a real CPU, and provides the same heatsink mechanical connection hole pattern as a CPU socket. Furthermore, its power is easily controlled, and its temperature monitored. The amount of heat generated by a real CPU cannot be precisely controlled, and its recorded temperature state is not completely reliable; thus, initial tests with the simulated CPU are better for describing the heat transfer characteristics of the tested cooling device. Moreover, these initial tests will uncover any potential problems with the cooling device, such as leaks, before it is tested in the much more expensive server system, which can cost in excess of \$10,000 CAD.

This section elucidates the design and construction of the simulated CPU. To this end, the formation of the heater itself, including the heated surface, case and insulation, is first described. This section subsequently covers the temperature sensor selection and placement, and, finally, the electrical power control and monitoring.

### 2.2.1 Heated Surface

The heated surface of the simulated CPU is made from a custom heating element and a copper block. The heating element uses double-twisted 28 AWG Nichrome wire, which was coiled and bent into a square, planar shape. The coiled wire sits in clay and forms a 40mm×40mm×5mm element. The overall electrical resistance of the heating element ( $R_{sim}$ ) is 20.2  $\Omega$ . To complete the heated surface, the heating element was pressed up against one of the large faces of a 40mm×40mm×5mm copper block such that the heat generated by the element is nominally transferred through the copper block. The other large face of the copper block acts as the surface being cooled, and it approximates the surface of a CPU's IHS.

### 2.2.2 Case and Insulation

A case, coupled with insulation, was designed and built to house and insulate the heated surface, as well as provide a connection interface for the CPU cooling device. The case was made with 28-gauge galvanized steel, machine screws and hex nuts. The overall shape of the case is that of a rectangular box (100mm×70mm×40mm) with one open side exposing the heated surface, which is set at the proper height using a steel stand inside the case. The unexposed sides of the heated surface are insulated in the case with fiberglass insulation (maximum temperature of 537 °C) and linings of silica fabric (maximum temperature of 982 °C) from McMaster-Carr. Lastly, the case was given flaps on two edges of the open side to make heatsink connection holes similar to a common CPU socket. (This socket is described in section 3.1.) The final case is shown in figure 2-5 with a connected cooling device.

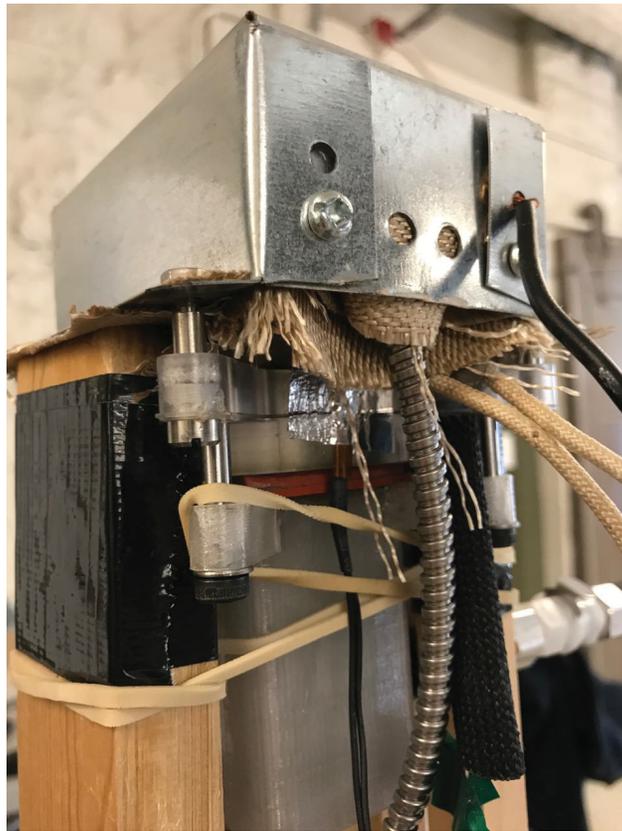


Figure 2-5: Simulated CPU with attached CPU cooling device.

### 2.2.3 Temperature Sensors

After the heated surface and case were finished, three temperature sensors were integrated into the simulated CPU. The first sensor was a 100  $\Omega$  platinum resistance temperature detector (RTD) probe (McMaster-Carr) that was placed inside the copper block, near its cooled surface, to give a reading of the approximate surface temperature ( $T_{surf}$ ). It has a temperature range from -20 °C to 176 °C and an accuracy of  $\pm 0.12\%$  of FS. Next, a heavy duty J type thermocouple probe (McMaster-Carr), with a range from 0 °C to 748 °C and an accuracy of  $\pm 0.75\%$  of FS, was inserted against the inside wall of the case to sense the approximate case temperature ( $T_{case}$ ). The final sensor added was an NTC thermistor surface probe (FrozenCPU), with a temperature range from 0 °C to 115 °C. This probe is attached to the most temperature-critical area of the tested cooling device with aluminum foil tape, and is monitored to ensure that the device does not overheat. This thermistor was calibrated using the same three-point thermistor calibration process and Steinhart-Hart equation (2.1) as described in section 2.1.6, which yielded coefficients:  $A = 1.02 \times 10^{-3}$ ,  $B = 2.32 \times 10^{-4}$ , and  $C = 2.62 \times 10^{-7}$ .

### 2.2.4 Electrical Power

The resistance-based heating element of the simulated CPU is powered by mains electricity. This section explains the electrical control system, which is important for the experiments, and the safety of the user and system that the power to the simulated CPU be controlled. The accuracy and some of the safety features integrated into the electrical power system of the simulated CPU are described below.

#### Power Control

The estimated maximum power required is about 600 W; therefore, the control system should be capable of varying the power to the simulated CPU from 0 W to (at least) 600 W. Moreover, it is desired that the entire experimental apparatus be digitally controllable during the tests. A digitally-controlled variac would suit these requirements, but a less expensive

alternative was selected instead. This alternative is a solid-state relay (SSR) that is used to control the average time the simulated CPU is powered on relative to the time it is powered off. The SSR selected works directly with mains power (120 VAC, 60 Hz). To avoid potential electrical problems in the high-voltage circuit, it is only capable of being turned on or off at the zero point (0 V) crossing of the AC sine wave (i.e. at a maximum frequency of 120 Hz). Since the power is in essence sampling the SSR at a frequency of 120 Hz, the SSR is operated at frequencies less than 60 Hz to avoid aliasing.

The resistance of the heating element is  $20.2 \Omega$  and the nominal input root mean square voltage ( $V_{RMS}$ ) is 120 V; therefore, the maximum average power ( $P_{ave}$ ) is calculated to be 713 W using:

$$P_{ave} = D \frac{V_{RMS}^2}{R_{sim}}, \quad (2.2)$$

where  $D$  is the duty cycle of the SSR, which is the fraction of time that the simulated CPU is in a powered state ( $D = 1$  at maximum power).

### **Power Sensor**

The voltage provided by the mains AC power drops slightly below its nominal 120 VAC under greater loads; therefore, it is important to monitor the voltage and current across the simulated CPU to accurately record the power. A digital power analyzer is ideal for this type of data acquisition; but, again, a less expensive option was selected. A P3IP4400 Kill A Watt Electricity Usage Monitor (abbreviated to KAW hereinafter) was modified (see figure 2-6) to send voltage and current readings to the control and data acquisition system. A normal KAW simply displays the voltage, current and power of the connected electrical device. This modified KAW was opened up to tap into the ground reference pin, and the 0-5 V current and voltage read pins of its LM2902 operational amplifier. Wires were soldered from these pins to external wire connectors drilled into the casing of the KAW, which were in turn connected to the control and data acquisition system.



Figure 2-6: Modified Kill A Watt Electricity Usage Monitor attached to custom breaker outlet.

### **Safety Features**

Since the heating element deals directly with mains power, multiple safety features were designed into the circuit of the simulated CPU to protect the user and the rest of the system. The first safety feature is at the power inlet. A custom electrical outlet breaker was created to permit remote shutdown of the simulated CPU power by the control and data acquisition system. This custom outlet consists of a power plug, an outlet box, an outlet, and a relay. The relay controls the line connection and can thus be used to turn the power to the outlet on or off. Also, the outlet used has a ground fault circuit interrupter that shuts down the power if it detects a loss of current to an unintended path, such as the ground. The KAW plugs into this outlet, and an extension cord, which connects to the two ends of the heating element, plugs into the KAW. The line connection of this extension cord is connected in series to two more components. One component is the SSR, and the other is a 7 A fuse

that protects against current surges. Finally, the ground line of the extension cord is used to ground the case of the simulated CPU.

## **2.3 Real CPU: Computer System and Construction**

In contrast to the simulated CPU experiments, the real CPU experiments are conducted using an actual commercially-available computer system. This section first summarizes the specifications of both the computer system and its CPU. Subsequently, the computer software used for testing is identified. Lastly, the construction of the computer stand and its assembly with the computer in the experimental apparatus is covered.

### **2.3.1 Computer System Specifications**

The system used in the real CPU tests is CIARA’s ORION HF320-G3 server (hereinafter referred to as the HF320). At the time of initial release (2015), this was CIARA’s highest performance system and generally considered the best in the industry. It was designed to be the fastest computer system in production to cater to companies such as high-frequency trading firms. Table 2–1 provides a summary of the server specifications, and figure 2–7 shows the inside of the server. The processor in this server is Intel’s (then) top-of-the-line CPU, known as the Intel® Core™ i7-5960X Processor Extreme Edition, which will be referred to as the 5960X herein. Table 2–2 provides specifications of the 5960X. To cool the CPU, the server employs an Asetek Liquid Cooling (5th generation) system, which (at the time) was considered the best processor cooling system on the market. The heat sink in their system uses a microchannel cooling method.

Table 2–1: CIARA’s ORION HF320-G3 server specifications.

ORION HF320-G3 Server	
Server Height	2U (3.50 inches or 88.9 mm)
Motherboard	ASUS X99-WS/IPMI Rev 1 with ASUS UEFI BIOS Utility - Advanced Mode - Version: 2.17.1246
Chipset	Intel® X99
CPU	Intel® Core™ i7-5960X Processor Extreme Edition (8 cores, 20 MB) - Serial Number: 2W517006B0055
Memory	4 × Corsair Dominator DDR4 8 GB 2666 MHz 1.35 V - Version: 4.23
Power Supply	770 W Redundant (1+1)
Cooling System	Asetek Liquid Cooling 5 <sup>th</sup> Generation
Fans	3 × SANYO DENKI Fans - Type: 9GA0812P1S61

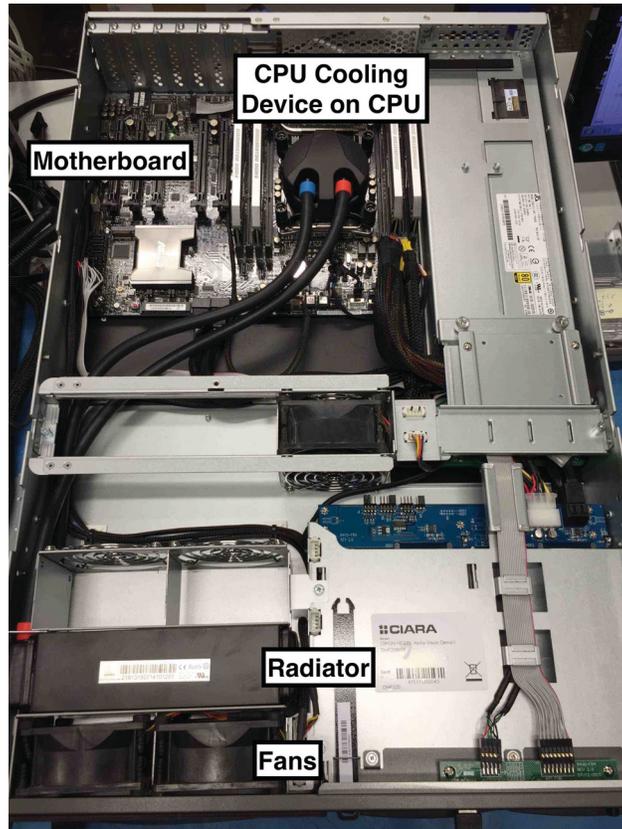


Figure 2–7: Inside CIARA’s ORION HF320-G3 server.

Table 2–2: Intel® Core™ i7-5960X Processor Extreme Edition specifications.

Intel® Core™ i7-5960X Processor Extreme Edition	
Lithography	22 nm
# of Cores	8
# of Threads	16
Cache	20 MB
Processor Base Frequency	3.00 GHz
Max Turbo Frequency	3.50 GHz
TDP	140 W
Sockets Supported	FCLGA2011-3

### 2.3.2 Computer System Software

The server is equipped with a suite of software that is used to test different overclocking settings. Upon booting the system, the overclocking settings can be set in the Unified Extensible Firmware Interface (UEFI), which is low-level software used to define the interface between the computer’s hardware and the operating system. The UEFI is the successor of the Basic Input-Output System (BIOS) and is the new standard for computer systems; however, manufacturers tend to refer to the UEFI as the UEFI BIOS to avoid potential confusion. The UEFI on the HF320 is ASUS UEFI BIOS Utility with Advanced Mode enabled. The graphical user interface (see figure 2–8) permits modification to numerous CPU settings, including the overclocking frequency.

Two programs are installed to the server’s Windows 7 operating system to conduct and record the overclocking tests: OCCT (version 4.2.0), which is an overclocking tool that can push the CPU utilization to 100% for a specified period of time; and HWiNFO64, which is a program that displays and records the data from all of the server’s embedded sensors. HWiNFO64 records, for example, the clock frequency and temperature of all eight CPU cores. Furthermore, it records the CPU package temperature, and the power to the CPU and server. However, the accuracy of this data is not known, so it is mainly used just for reference. The windows of both applications are shown in figure 2–9.



Figure 2-8: ASUS UEFI BIOS Utility graphical user interface.

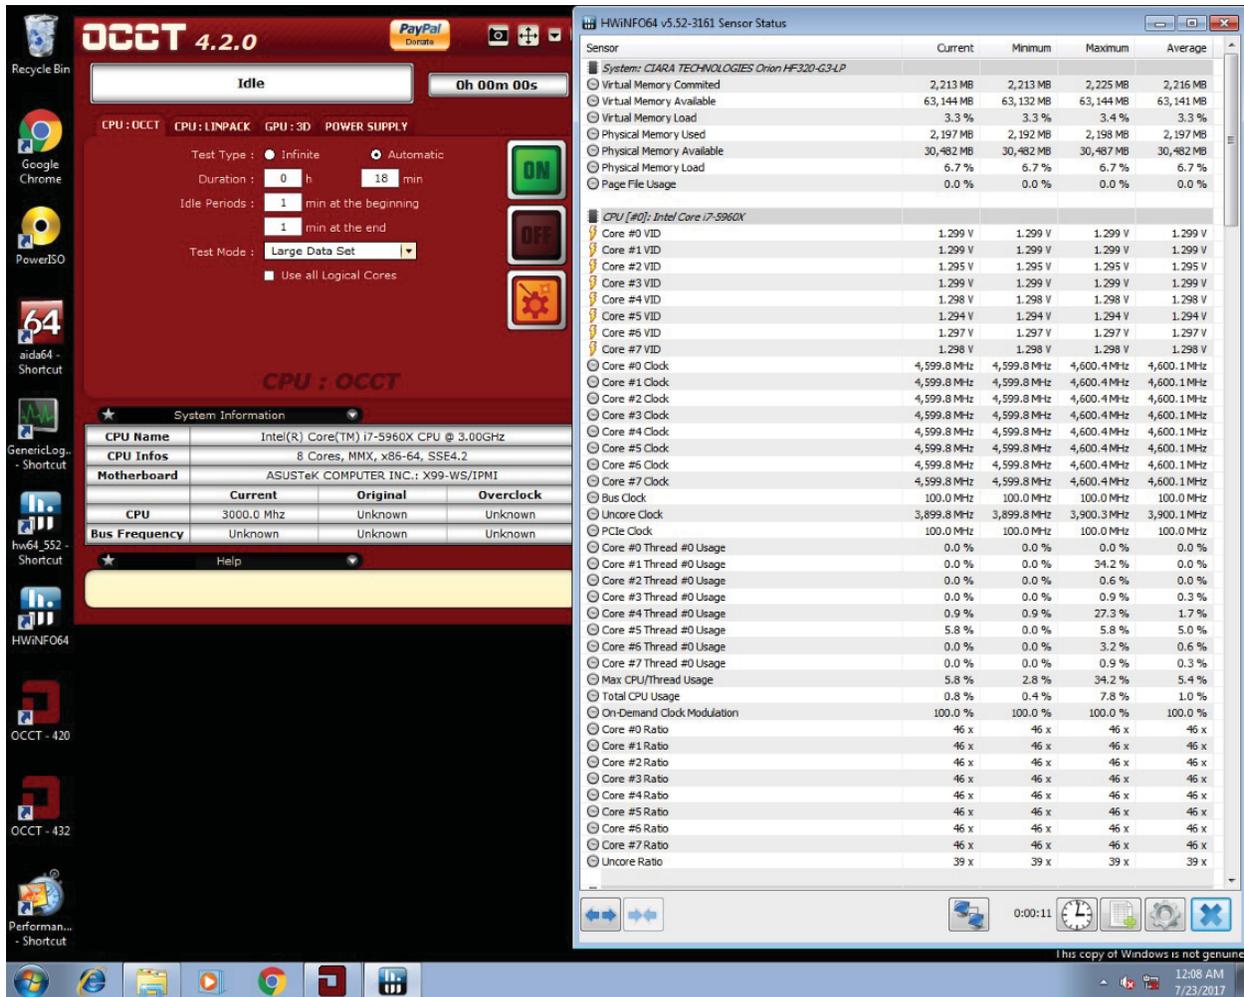


Figure 2-9: OCCT and HWiNFO64 application windows.

### 2.3.3 Stand Construction and Assembly

A stand was built for the server to satisfy space limitations, and add the ability to change the height of the server and place it upside down. The stand was built out of UNISTRUT®, which is easily customizable metal framing material, and it was bolted into the wooden platform and clamped to the table top above the electronics system of the experimental apparatus. Figure 2–10 shows the assembled server and stand with the remainder of the experimental apparatus.

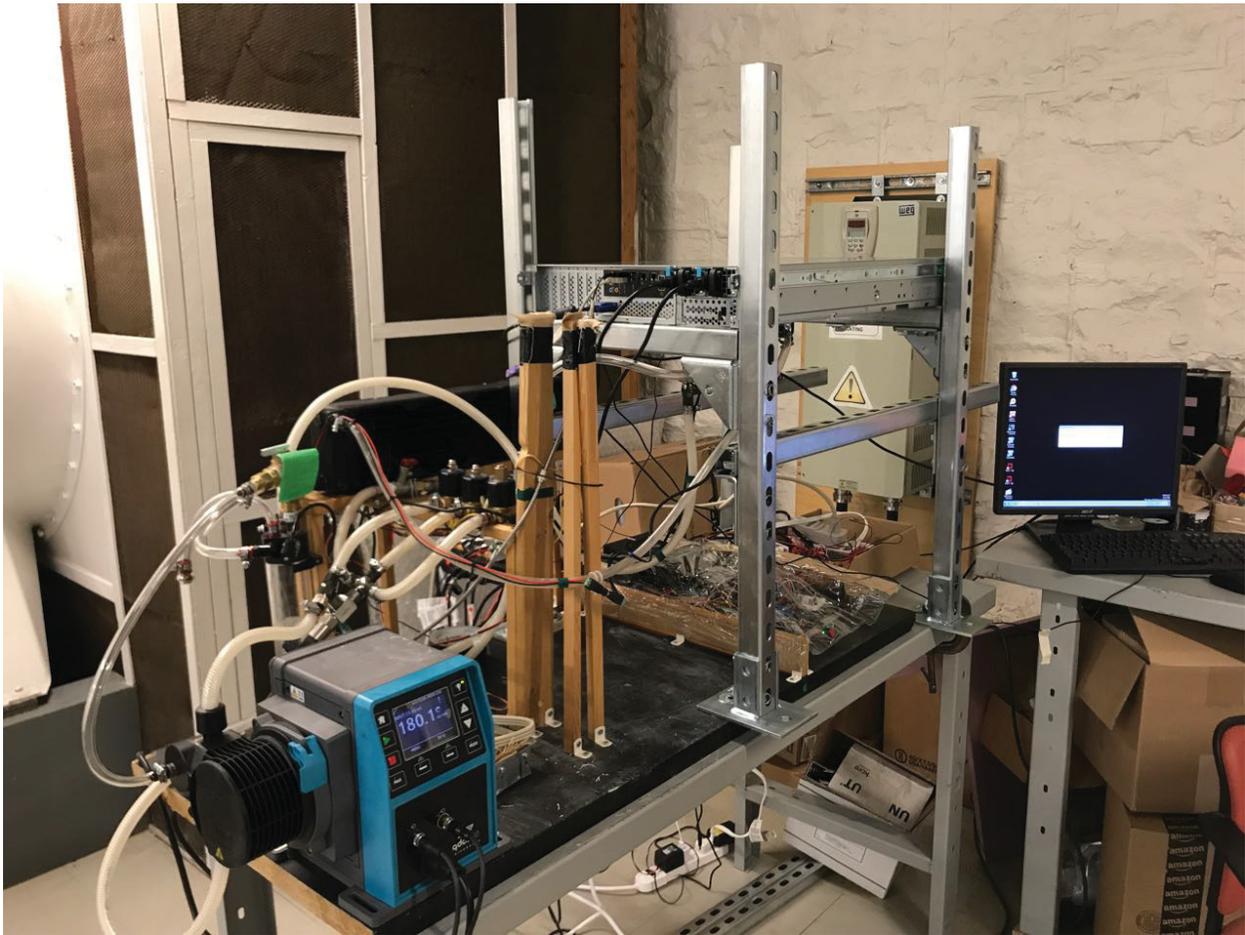


Figure 2–10: Experimental apparatus with server in stand.

## 2.4 Control and Data Acquisition

This final section of the experimental apparatus chapter is focused on all of the components, connections and programming that went into building the control and data acquisition

system. First, the master controller's selection is justified. Then, the controllers that are slave to the master controller are identified. Subsequently, an overview of the input and output connections is given. Finally, the application programmed to operate the control and data acquisition system is described.

#### **2.4.1 Master Controller: Raspberry Pi**

A Raspberry Pi 3 Model B single-board computer (abbreviated as RPi hereinafter) was selected as the master controller for the entire control and data acquisition system. Despite being a full-fledged computer, the RPi is the size of a credit card, which greatly increases the portability of the experimental apparatus. It is also inexpensive (50 CAD) and supported by a large, experienced community. Furthermore, it is equipped with 40 general-purpose input/output (GPIO) pins that enable it to connect to numerous devices. A few more relevant specifications are highlighted below:

- quad-core 1.2 GHz Broadcom BCM2837 64 bit CPU;
- 1 GB of RAM;
- ethernet, wireless and bluetooth connectivity;
- full-sized HDMI port;
- 4 × USB 2 ports;
- micro SD port; and
- micro USB port for power (up to 2.5 A).

The micro SD port contains a micro SD card that loads the operating system and stores data. Raspbian is the operating system used by the RPi, and it is based on Debian Jessie, which uses the Linux kernel. This operating system runs a full desktop environment and offers most of the functionality of other Linux-based systems. The most commonly used programming language on the RPi is Python.

The 40 GPIO pins and other ports of the RPi are all easily accessible and controllable using Python, which allows powerful programs to be built directly on the RPi. A few of the pins are dedicated ground or power (3.3 or 5 V) pins, but the remainder are GPIO and are

capable of digital input/output (0/3.3 V) signals, PWM, and reading interrupts. Moreover, many of the pins have alternative functions that allow the RPi to communicate with other devices using multiple protocols, including UART (Serial), I<sup>2</sup>C and SPI.

### 2.4.2 Slave Controllers

The master controller alone is not sufficient to control the entire system. It is supported by two additional controllers. The purpose of the first slave controller is to increase the functionality of the RPi. The second is a controller for some of the components of the simulated CPU.

#### **Additional Inputs and Outputs: Arduino Mega**

Although the RPi has 40 connection pins and four USB ports, it was determined that it alone would not be capable of supporting all of the required connections for the control and data acquisition system; therefore, an Arduino Mega (abbreviated to Mega) microcontroller board, based on the ATmega1280, was added to the system. This board adds another 54 digital pins (0/5 V) that have the same capabilities as the GPIO pins of the RPi, including the communication protocols. In addition to the digital pins, the Mega also has 16 analog pins that are capable of reading analog voltage signals from 0 to 5 V. These pins convert the analog signal to a digital signal with 10 bits of resolution. The RPi does not contain any analog pins; therefore, this addition greatly expands the capabilities of the system.

The Mega communicates with the RPi via a UART connection using a USB cable. The Mega is not a stand-alone computer like the RPi, but a program written in Arduino Programming Language can be uploaded to it to run on a loop. A program known as nanpy (developed by Andrea Stagi and released under MIT License) was uploaded to the Mega, and it allows the RPi to control the Mega as a slave over UART using Python with the nanpy library. This simplified the overall program development process.

### **Simulated CPU: Arduino Nano**

An additional Arduino was added to the system to handle the more electrically-dangerous components of the simulated CPU. It is an Arduino Nano (abbreviated to Nano), which is a less powerful and less expensive version of the Mega with fewer connection pins. The simulated CPU temperature, voltage and current readings are all handled by the Nano. The temperature sensor connections are considered potentially dangerous because, in the case of a malfunction, they could come in contact with the mains power lines of the heating element. Furthermore, the modified nature of the KAW makes the voltage and current connections inherently dangerous. The outlet breaker circuit connection is safely handled by the RPi because the connection to the relay board is isolated from the solenoid relays with an optocoupler. The SSR connection is similarly isolated.

To fully electrically isolate the simulated CPU circuit from the main circuitry, a protected communication connection between the Nano and the RPi was implemented. Again, the communication protocol used is UART; however, a USB cable could not be used to make the connection in this case. Instead, two optocoupled circuits were created to connect the TX and RX pins on both boards using optocouplers (6N137M) capable of fast enough data transmission.

The voltage and current pins tapped into in the KAW output the instantaneous voltage and current values rather than root mean square (RMS) values. Therefore, the controllers must sample at a high enough frequency to generate an accurate measurement of the voltage and current RMS values. Unfortunately, the nanpy program used with the Mega is not capable of sampling at the required rate; thus, a custom Arduino program was created to handle the high-frequency sampling and RMS calculations. Upon receiving the proper signal code from the RPi over the UART Serial connection, the program on the Nano takes 256 samples of the voltage and current simultaneously at a sample frequency of about 1,950 Hz. Then, it calculates the average power, and the voltage and current RMS values from those

samples using the following equations:

$$P_{ave} = \frac{\sum_1^{256} v_j i_j}{256}, \quad (2.3)$$

$$V_{RMS} = \sqrt{\frac{\sum_1^{256} v_j^2}{256}}, \text{ and} \quad (2.4)$$

$$I_{RMS} = \sqrt{\frac{\sum_1^{256} i_j^2}{256}}, \quad (2.5)$$

where  $I_{RMS}$  is the RMS current of the simulated CPU, and  $v_j$  and  $i_j$  are the instantaneous voltage and current, respectively, at sample point  $j$ . Finally, the average power, RMS voltage, RMS current, as well as the temperatures for the three simulated CPU sensors are all sent to the RPi over the serial connection.

### 2.4.3 Inputs

The controllers of the control and data acquisition system handle a number of input connections. These connections transmit various types of signals that must be conditioned in different ways to satisfy the input requirements of the controllers. This conditioning is achieved with custom circuits that, unfortunately, cannot be described in the present thesis due to length restrictions. The input circuits consist of mainly basic elements such as resistors and transistors. Furthermore, Zener diodes are used to protect many of the input pins from overvoltage. The maximum input voltages for the RPi and the Arduinos are 3.3 V and 5 V, respectively.

Table 2–3 summarizes the input connections, along with their connected controller, connection output signal and controller input signal. The RPi generally handles the on/off type connections, the digital signals and the interrupt signals. The Arduinos utilize their ability to read analog signals. Therefore, the Mega is connected to all of the temperature and pressure sensors, and the Nano handles the remaining analog connections deemed too dangerous for direct connection to the main circuitry.

Table 2-3: Input connections and their signal types.

Input Connection	Controller	Connection	Output Signal	Controller Input Signal
Pump - Run Status	RPi	Open Collector		0/3.3 V
Pump - Alarm	RPi	Open Collector		0/3.3 V
Main Reservoir - Liquid Level Alarm	RPi	Magnetic Reed Close Contact	Open/ Close Contact	0/3.3 V
Environment - Ambient Temperature	RPi	Digital (DHT22-AM2302)		Digital (DHT22-AM2302)
Environment - Ambient Humidity	RPi	Digital (DHT22-AM2302)		Digital (DHT22-AM2302)
Fans - Tachometer	RPi	Open Collector		Interrupt
Pump - Flow Rate	Mega	4-20 mA		0-5 V
Main Reservoir Outlet - Temperature	Mega	Resistance (10 k $\Omega$ Nom.)		0-5 V
CPU Cooling Device Inlet - Temperature	Mega	Resistance (10 k $\Omega$ Nom.)		0-5 V
CPU Cooling Device Outlet - Temperature	Mega	Resistance (10 k $\Omega$ Nom.)		0-5 V
Radiator Inlet - Temperature	Mega	Resistance (10 k $\Omega$ Nom.)		0-5 V
Radiator Outlet - Temperature	Mega	Resistance (10 k $\Omega$ Nom.)		0-5 V
Main Reservoir Inlet - Pressure	Mega	4-20 mA		0-5 V
Pump Outlet - Pressure	Mega	4-20 mA		0-5 V
CPU Cooling Device Outlet - Pressure	Mega	4-20 mA		0-5 V
Simulated CPU - Voltage	Nano	0-5 V		0-5 V
Simulated CPU - Current	Nano	0-5 V		0-5 V
Simulated CPU Surface - Temperature	Nano	4-20 mA		0-5 V
Simulated CPU Case - Temperature	Nano	4-20 mA		0-5 V
CPU Cooling Device - Temperature	Nano	Resistance (10 k $\Omega$ Nom.)		0-5 V

### 2.4.4 Outputs

Unlike the input connections, there is no benefit to connecting the outputs to the slave Arduino controllers; therefore, all output connections are connected directly to the master controller. The output connections are listed in table 2–4, including their controller output signal and connection input signal. Like the input connections, signal modification is necessary to meet the signal requirements, but, in this case, the circuits are slightly more complex. Many of the circuits require relays and external power supplies, and one circuit, namely the pump flow rate signal conditioning circuit, needs an SPI to 4–20 mA 12-bit digital-to-analog converter (MCP4921).

Table 2–4: Output connections and their signal types.

Output Connection	Controller	Controller Output Signal	Connection Input Signal
Pump - Run/Stop	RPi	0/3.3 V	0/5 V
Pump - Flow Rate	RPi	SPI	4–20 mA
Fans - Speed	RPi	PWM	PWM
CPU Cooling Device - Normally Open Valve	RPi	0/3.3 V	0/12 V
Main Reservoir - Normally Closed Valve	RPi	0/3.3 V	0/12 V
Backup Reservoir - Normally Closed Valve	RPi	0/3.3 V	0/12 V
Alarm	RPi	0/3.3 V	0/3.3 V
Simulated CPU - Outlet Breaker	RPi	0/3.3 V	0/3.3 V
Simulated CPU - Power	RPi	PWM	PWM

### 2.4.5 Application

With the help of the nanpy program on the Arduino Mega and the custom Arduino program on the Nano, the RPi master controller has control over all of the electronic input and output connections. Consequently, a program application was written in Python on the RPi to handle all of the control and data acquisition of these connections. This program has five major components: a library of connection drivers, a controller, a data recorder, a safety monitor, and a graphical user interface. This section covers each of these components.

The library of connection drivers can be split into two main categories: input connection and output connection drivers. The purpose of the input connection drivers is to convert the digital representations of the electrical signals read by the microcontroller pins to the sensor values they represent. For example, the digital 10-bit value read from one of the thermistor temperature sensor connections is first converted back to its analog voltage representation (0–5 V); then, that voltage is translated into the resistance of the sensor using its circuit model, and, finally, that resistance is used to calculate the temperature value using equation 2.1 and the Steinhart-Hart coefficients appropriate for that sensor. The output connection drivers perform a similar function but in reverse. The driver must determine what digital signal to send to the output pin to achieve the desired outcome from the output connection. In most cases, the driver must simply know whether to drive the pin to a low or high state (0/3.3 V) since most of the output connections have only two potential functions: on or off.

The controller manages all of the connections, which all have associated connection drivers. Once every second, it updates the value of each connection. Furthermore, it permits control of the output connections. Lastly, it handles the startup, initiation of the connections, and the shutdown procedure, which typically involves ensuring the outputs are in their proper state.

All connection values are updated by the controller, but it is the data recorder that keeps track of those values over time. Each connection’s value is recorded every second and written to a single data file along with a time stamp. Moreover, the data recorder also records log items along with their time stamps. Log items are automatically generated messages designed to give insight into what is going on with the experimental apparatus and notify the user of any problems.

The safety monitor also tracks connection values, but instead of recording them, the monitor ensures that they are within certain safety limits. Every update cycle (1 Hz), the safety monitor goes through a series of safety cases to check for any potential safety issues or problems. Each safety case typically checks one critical connection value. Some of the safety

cases are as simple as checking if the pump or the main reservoir liquid level alarms are on. If these alarms are triggered, the safety monitor uses the controller to stop the experiment (turn off the simulated or real CPU, and the pump), to set the fans to maximum speed, and to open the main reservoir valve to send liquid coolant back to the reservoir. Moreover, a system alarm is sounded, and a log item explaining the problem is created. The safety cases for the critical analog connections (such as the temperature, pressure, voltage, current and power sensors) are more complex because they generally split the safety actions into multiple levels that each generate a safety warning log item at a minimum. For example, the pressure in the main reservoir must be limited to 2 bar (gauge); therefore, a safety case exists for the pressure sensor in the main reservoir with three different levels. The first level is set at 1 bar, and if the pressure exceeds that, then the system alarm sounds. If the pressure continues to increase and reaches the second level limit (1.3 bar), then the experiment is stopped, and if it reaches level three (1.6 bar), then both reservoir valves are opened to release the pressure to the much larger backup reservoir. Therefore, the mechanical 2 bar safety valve should only ever become necessary if there is a system malfunction. Furthermore, there is also a zeroth level that ensures the pressure doesn't drop below -0.2 bar, which would indicate a major sensor malfunction. In this case, the system would commence a safe shutdown procedure.

The graphical user interface of the application consists of a single large window with five panels (see figure 2-11). The large panel on the left displays all of the input and output connections, as well as their real-time values and units. The largest panel on the right is responsible for showing the log items in real time, along with their time stamps. Next, the bar along the bottom of the window is the status bar, which also indicates the name of the data file that is being created during the experiment. The status bar colour turns red when an alarm is sounded. Towards the top, there is a control button bar that allows the user to start the experiment, open valves, change the pump flow rate, etc. The time elapsed since the application was launched is also displayed on this bar. Finally, at the very top, is the menu bar. Currently, this menu has only one dropdown menu that displays a number of

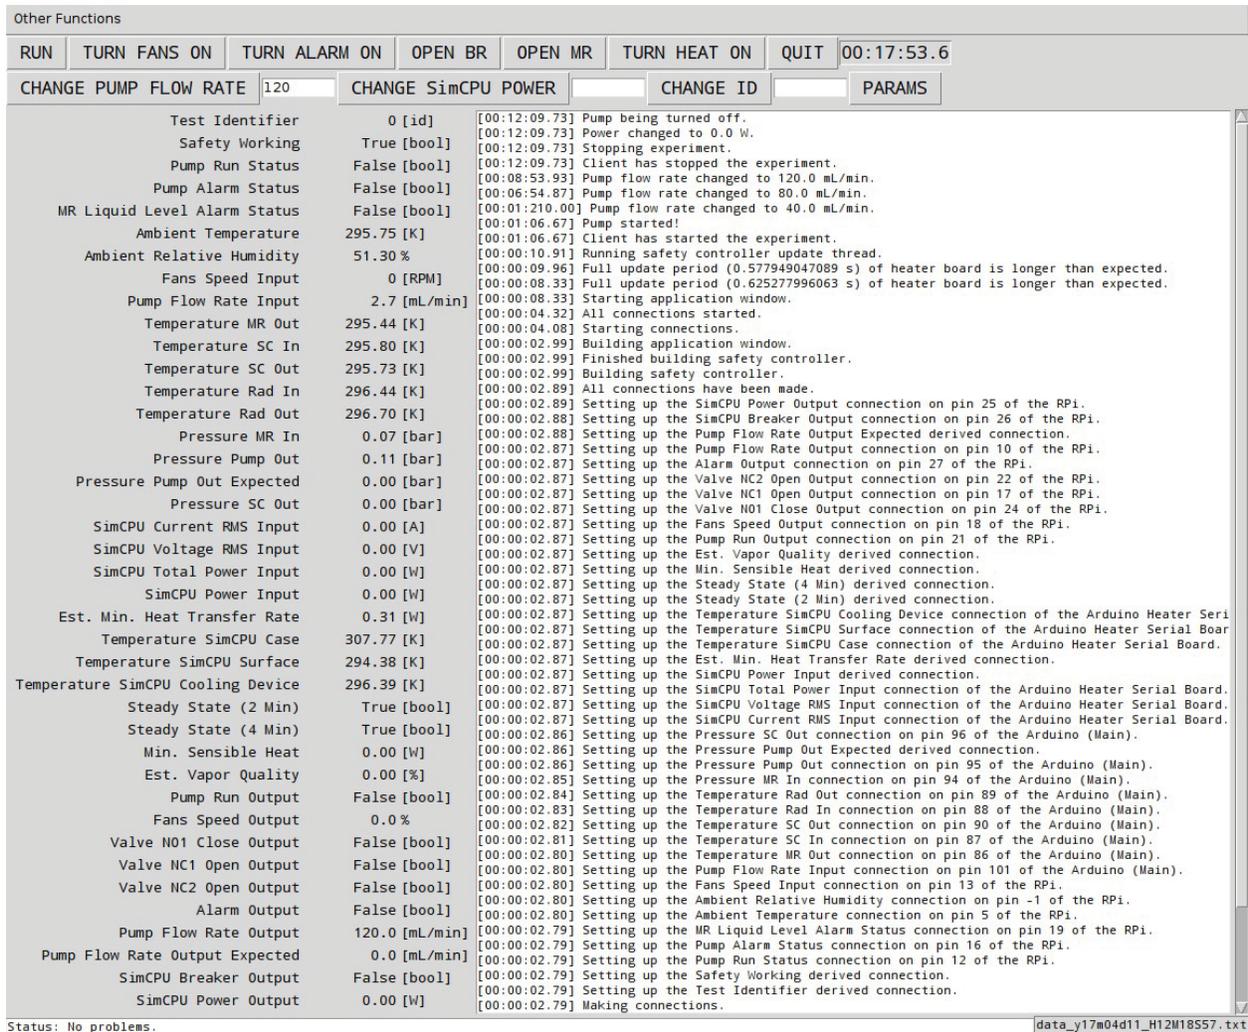


Figure 2–11: Window of the application’s graphical user interface.

functions, mainly concerned with initial calibrations, that were not necessarily appropriate for the control button bar.

A closer inspection of the left panel displaying the connection values reveals that there are actually a number of values listed that do not correspond to any of the input or output connections. These false connections are used to display important information in real time alongside the real connection values. A number of these connections show a value that is derived from the values of other connections. For example, a derived value for the expected pump outlet pressure based on the pump flow rate and the flow characteristics of the current cooling device being tested is next to the value of the actual pump outlet pressure. Thus,

the two can be compared, and if there is a large discrepancy, then the user would be alerted to the problem. Additionally, there are two steady state connections that monitor critical connection values, such as the simulated CPU's surface temperature, over two different time scales (2 and 4 minutes) to determine whether the system is in a steady enough state to take a measurement. Once a measurement is ready to be taken, another false connection comes into use. The test identifier is simply an integer that is used to indicate which test run is currently underway. The test identifier is normally set to 0, and when a measurement is ready to start, the identifier is set to the proper test run identifying integer. After the measurement is complete, the test identifier is set back to 0. This feature is extremely useful for data processing after the experiment.

## **CHAPTER 3**

### **CPU Cooling Device**

The following chapter describes the CPU cooling device designed and prototyped to be used in the validation tests of the experimental platform. This cooling device is also an initial attempt at improving processor cooling and achieving enhanced processor performance. First, the objectives and requirements for the design of the device are stated. Next, the cooling technology selected and the conceptual design of the cooling device are described. Then, the design of an initial prototype that approximates the conceptual design is detailed, and the resulting product is shown.

#### **3.1 Design Objectives and Requirements**

The main objective of this processor cooling design is to incorporate a high heat flux cooling technology with the potential of improving processor performance into a novel CPU cooling device. CIARA's ORION HF320-G3 server is capable of operating at an overclocked processor speed of 4.7 GHz with its current cooling system; therefore, achieving reliable performance of this server with processor speeds above 4.7 GHz is a major goal. This is to be accomplished by designing a device that is capable of absorbing more heat than the current device in the HF320 while maintaining processor core temperatures below their critical level. In the case of the Intel® Core™ i7-5960X Processor, the critical core temperature is around 85 °C, and Intel suggests that the maximum IHS temperature be 66.8 °C.

In addition to the above, the design must also be compatible with existing computer systems, such as the HF320, and not require any significant modifications to the systems. Particularly, the design should not require any modification to the CPU because this voids the CPU's warranty. Moreover, it is beneficial to have a device of the smallest possible size.

The device should furthermore be safe for the computer system, the environment, and any potential handlers. Specifically, leaks are a common problem for liquid cooling solutions that must be avoided. Additionally, the device should not produce excessive vibrations or noise, which can ultimately reduce its reliability.

The overall cost of the device, including manufacturing, assembly and operating costs, should be justified by the relative cooling performance so that the device could justifiably be used in the real world. The manufacturing and assembly costs are associated with material costs and complexity, while the operating costs are associated with the device efficiency, which, in this case, is defined as the amount of heat removed per unit of electrical power necessary to run the device.

For testing purposes, the proof-of-concept prototype of the cooling device is required to be compatible with the 5960X and its LGA 2011-3 socket in particular. The processor socket is coupled with an independent loading mechanism (ILM) that serves to properly seat the processor in the socket as well as provide a connection interface for the heat sink. Figures 3-1 and 3-2 show how the processor, socket and ILM come together.

The processor’s cooling device mechanically connects to the ILM via four heat sink M4 studs at the outer corners. The centers of these four studs are spaced such that they form an 80 mm by 80 mm square. The relevant vertical dimensions are given in figure 3-3 and table 3-1.

Table 3-1: Intel® LGA-2011-3 socket vertical dimensions reproduced from Intel® Core™ i7 Processor Family for the LGA2011-3 Socket - Thermal/Mechanical Specification and Design Guide (TMSDG) (2014).

Dim A - Motherboard to top of heat sink stud	4.678(+0.367/-0.231) mm
Dim B - Motherboard to top of IHS load lip	6.581±0.289 mm
Dim C - Motherboard to top of IHS	8.481±0.279 mm

Finally, the exposed cooling area of the IHS and its dimensions are shown in figure 3-4. The IHS is made of a nickel-plated copper.

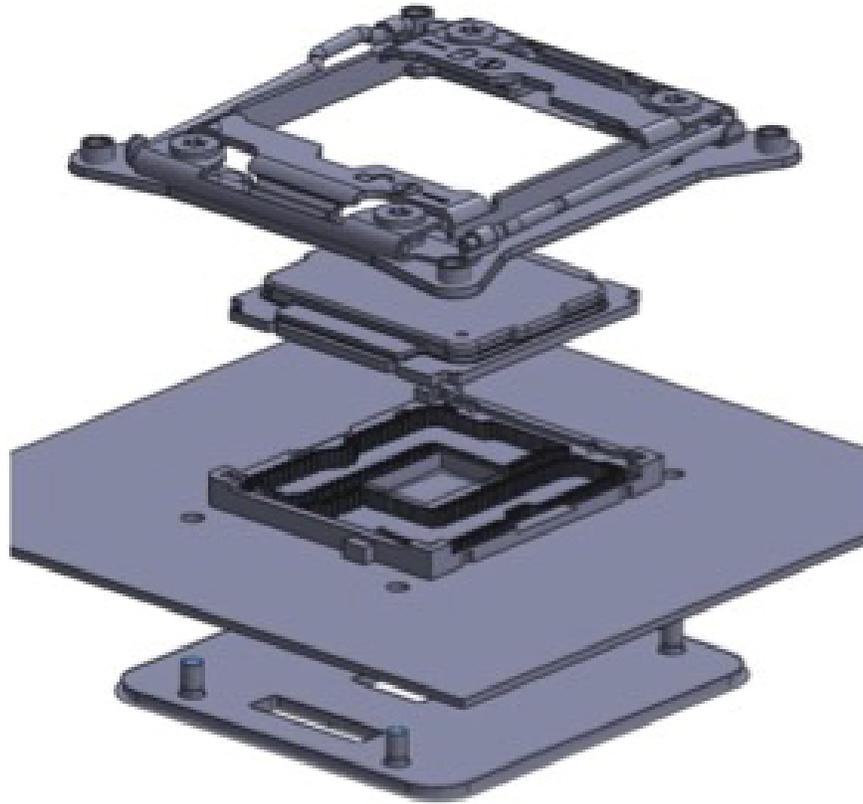


Figure 3-1: Intel® LGA-2011-3 socket exploded view reproduced from Intel® Core™ i7 Processor Family for the LGA2011-3 Socket - Thermal/Mechanical Specification and Design Guide (TMSDG) (2014).

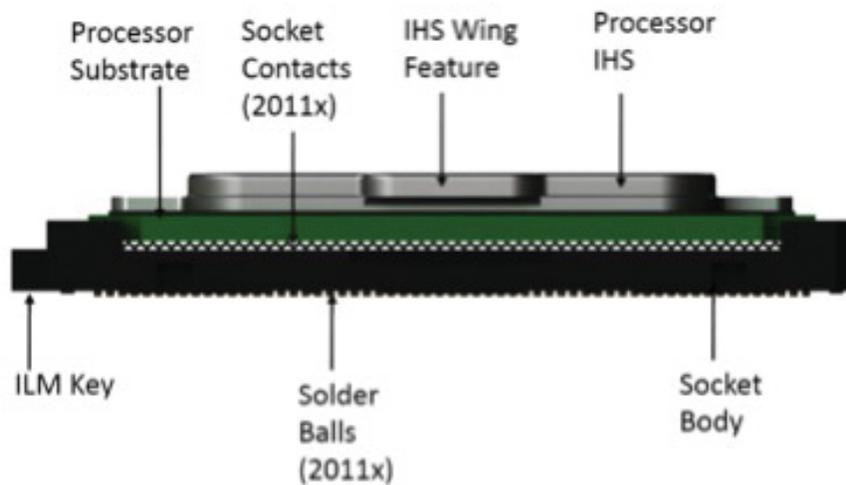


Figure 3-2: Intel® LGA-2011-3 socket labeled view reproduced from Intel® Core™ i7 Processor Family for the LGA2011-3 Socket - Thermal/Mechanical Specification and Design Guide (TMSDG) (2014).

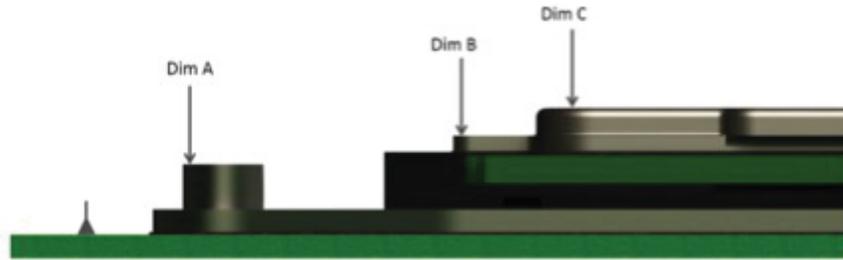


Figure 3–3: Intel® LGA-2011-3 socket vertical dimensions reproduced from Intel® Core™ i7 Processor Family for the LGA2011-3 Socket - Thermal/Mechanical Specification and Design Guide (TMSDG) (2014).

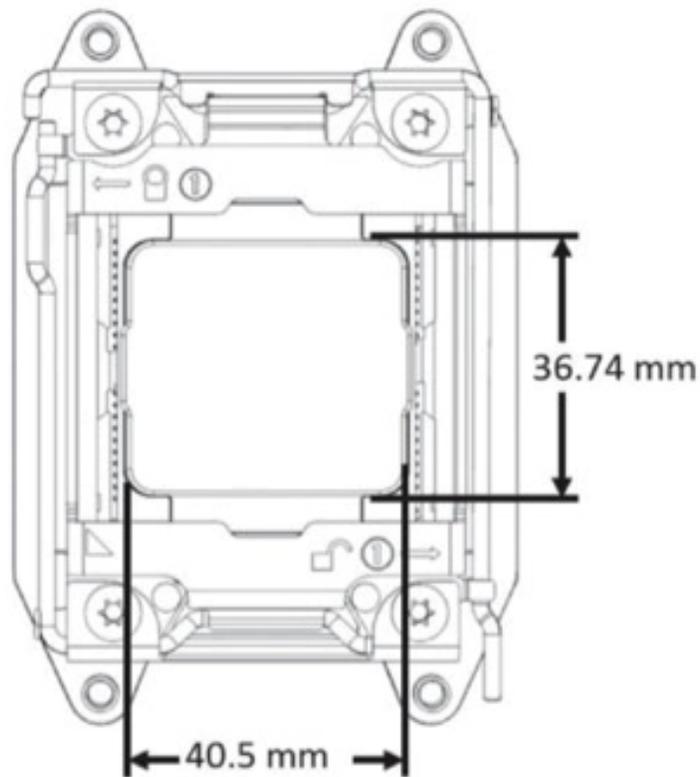


Figure 3–4: Intel® Core™ i7 Processor IHS dimensions reproduced from Intel® Core™ i7 Processor Family for the LGA2011-3 Socket - Thermal/Mechanical Specification and Design Guide (TMSDG) (2014).

### 3.2 Conceptual Design Overview

With an understanding of the design objectives and requirements for the CPU cooling device, an overview of the final conceptual design can be given. This overview starts with

the selection of the cooling technology and justification for its selection, and then proceeds to describe how that cooling technology is leveraged in the design.

After considering the various high heat flux cooling methods for flat surfaces during the literature review, spray cooling was selected as the method with the most potential to achieve an improvement in processor cooling. Spray cooling offers numerous advantages. It is well-suited for two-phase cooling, and therefore can utilize a coolant's latent heat of vaporization, which is capable of absorbing a large amount of heat. Other cooling methods, such as microchannels, suffer from significant heat transfer non-uniformity, especially when used in a two-phase fashion, and this is a problem for processors, which operate less reliably when subjected to large heat transfer and temperature gradients. Spray cooling, however, avoids this problem because heat transfer, as well as temperature, is characteristically uniform across the spray surface (Bostanci *et al.* 2009). Furthermore, spray cooling can be operated effectively at very low flow rates, and multiple heat transfer improvements are possible in the form of spray surface enhancement (Pais 1992; Silk 2006).

However, the biggest potential advantage of spray cooling is the ability to target the heated surface directly. Traditional liquid cooling systems must employ a cold plate and a TIM between the IHS and the coolant. By targeting the IHS directly, two thermal layers (the TIM and cold plate) and two thermal contact resistances (between the IHS and TIM, and between the TIM and cold plate) are instantly removed. This greatly reduces the overall thermal resistance of the cooling system.

These strengths of spray cooling are fully leveraged in the conceptual design of the CPU cooling device, which specifically takes full advantage of direct spray impingement on the IHS. This conceptual design is shown in figure 3–5. Essentially, the device is a spray chamber with the IHS being the bottom surface that is sprayed. The spray chamber (10) is split up into two pieces: a spray chamber wall (11) and a spray chamber lid (12). The two pieces connect via a fastening mechanism (16) that also creates a hermetic seal between them. Arms (15) extend from the spray chamber wall and fasten to the heat sink studs with four

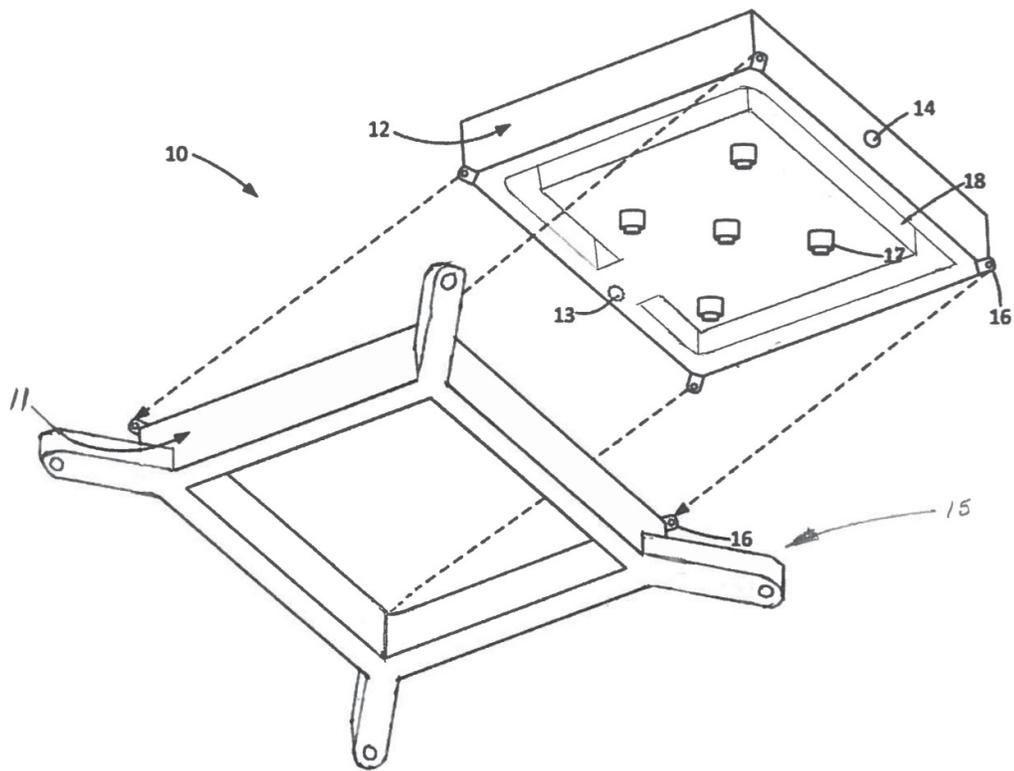


Figure 3–5: CPU spray cooling device conceptual design reproduced from U.S. Patent File No. P3711PC00 (Hinton and Mydlarski 2016).

M4 screws. The spray chamber wall forms a hermetic seal along the outer perimeter of the surface of the IHS; therefore, the entire chamber is hermetically sealed onto the surface of the IHS. During operation, coolant flows through the inlet (13) of the spray chamber lid and is forced through an array of spray nozzles (17). The coolant atomizes as it passes through the spray nozzles and subsequently impacts and cools the IHS. The heated coolant then collects in the inner perimeter of the spray chamber lid (18) and flows through the outlet (14).

Although a single-piece spray chamber design is certainly possible, the two-piece spray chamber is the preferred design embodiment. This feature enables safe removal of coolant off of the IHS upon removal of the spray chamber lid without the chance of a leak occurring. Furthermore, it facilitates maintenance and cleaning of the parts, and it allows easy customization.

Excess liquid cannot be allowed to pool on the IHS because this significantly reduces heat transfer. The current design does not deal with this problem in a sophisticated fashion due to certain constraints. One constraint is that an overflow system cannot be added around the perimeter of the IHS because the IHS is practically flush with the motherboard. The current strategy to prevent pooling is simply to operate the spray system in an inverted manner, such that the excess coolant flows downward toward the outlet.

### **3.3 Prototype Detailed Design**

The conceptual design above represents the desired end product. However, a less constrained and less complex design was prototyped first to reach a testable device more quickly. This prototype was thus not restricted in size and was given a single nozzle rather than an array of nozzles. This section focuses on the main design aspects of the CPU cooling device prototype. The coolant and spray nozzle selections follow, and then a description of the spray chamber design is given.

#### **3.3.1 Coolant**

For the proof-of-concept testing, two coolants were tested: one single-phase coolant and one two-phase coolant. For single-phase coolants, the coolant's sensible heat is the most important property because higher sensible heat means more energy per degree Celsius is absorbed per unit mass of coolant. For two-phase coolants, the latent heat of vaporization and the standard boiling point (saturation temperature at 1 bar) are the most important characteristics. The latent heat of vaporization should be high so that the coolant is capable of absorbing large quantities of heat when it vaporizes on the surface of the IHS. The boiling point must also be below the maximum allowable temperature of the surface of the IHS (66.8 °C in this case) so that the coolant vaporizes when impinging on the IHS. Also, it should be higher than the ambient temperature such that the coolant is capable of condensing to liquid form when it releases heat to the environment. Ambient temperature is nominally

20 °C, but it can be as high as 30 °C locally. Therefore, the two-phase coolant should have a boiling point in the range of 30 °C to 66.8 °C.

Other considerations when selecting a spray coolant are price, viscosity, surface tension, dielectric strength, material compatibility, safety, and environmental considerations. Low viscosity coolants are preferred because a highly viscous coolant is difficult to pump and spray. High surface tension can be desirable to decrease the chance of leaks around fittings; however, higher surface tensions also decrease the surface wettability of the coolant, which can decrease heat transfer. A high dielectric strength coolant is desirable, especially for initial testing, because electronic components in the server will not become damaged in the event of a leak. In terms of material compatibility, the coolant must not corrode or react with any component in the fluid loop. Thus, inert coolants are best. Safety considerations such as toxicity, flammability, and explosive risk must all be taken into account in order to protect anyone who could come into contact with or be in the vicinity of the coolant. Finally, environmental concerns include ozone depletion potential and global warming potential.

The single-phase coolant selected for initial testing is distilled water, and the two-phase coolant selected is 3M's Fluorinert™ Electronic Liquid FC-3284, which is abbreviated to FC-3284 in this thesis. The most relevant properties for these two coolants are in table 3–2. Properties are stated at 25 °C and 1 bar except for the specific heat capacities because the average specific heat capacities over the expected temperature range is more useful for future calculations. The specific heat capacity of water is the average of the value at 20 °C and the value at the maximum IHS surface temperature (66.8 °C). The specific heat capacity of FC-3284 is the average of the value at 20 °C and the value at its saturation temperature (50 °C). Surface tension is measured in contact with air.

Water is the obvious choice for the single-phase coolant. It's thermal properties are unparalleled, and it is abundant and cheap. Using distilled water has the added benefits of being dielectric and reducing potential contaminants. Furthermore, water has fairly good

Table 3–2: Relevant properties of distilled water and 3M’s Fluorinert™ Electronic Liquid FC-3284 taken from NIST (2016) and 3M Fluorinert™ Electronic Liquid FC-3284 Product Information (2000).

Coolant	Specific Heat $c_{p,l}$ [kJ/kg-°C]	Latent Heat $\lambda$ [kJ/kg]	Boiling Point $T_{sat}$ [°C]	Liquid Density $\rho_l$ [kg/m <sup>3</sup> ]	Dynamic Viscosity $\mu_l$ [cP]	Surface Tension $\sigma_l$ [N/m]
Distilled Water	4.1847	2230	99.974	997.06	0.89008	0.07199
FC-3284	1.068	105	50	1710	0.72	0.0130

material compatibility, and it is safe both to handlers and to the environment. If the saturation temperature of water was in the necessary range, then water would be ideal for two-phase spray cooling as well. It is true that its saturation temperature can be lowered by reducing the system pressure; however, that is not tested in the current study due to the added complexity and cost.

Although FC-3284 has a considerably lower specific heat capacity and latent heat of vaporization compared to water, it is the best coolant available with a boiling point in the necessary range. The boiling point also sits in the middle of that range, which is advantageous to both the vaporization and condensation processes. Furthermore, it has a high dielectric strength, is safe, and has good material compatibility. FC-3284 has zero ozone depletion potential, but it has fairly high global warming potential so it is only safe to use in a closed environment. The only other disadvantage of the coolant is that it is expensive (180 CAD/litre).

### 3.3.2 Nozzle

For proof-of-concept testing, a single spray nozzle is used in the prototype design rather than an array of multiple nozzles. The idea is to use a single high-quality nozzle to prove that spray cooling has the necessary heat transfer abilities first before dealing with the added complexity of multiple nozzles. A nozzle with the capability of delivering a full spray at very low flow rates is desired. A low flow rate capability allows the minimum flow rate required

at different CPU powers to be determined. Furthermore, the capability to function at low flow rates allows the spray system to be potentially used in a non-inverted fashion when using the two-phase coolant because the flow rate can be set low enough so that it always completely vaporizes.

Multiple spray nozzle types exist, including full cone, hollow cone, misting and fan. Hollow cone and fan nozzles can achieve very low flow rates; however, their spray areas are not ideal for this application. Hollow cones spray the circumference of a circle, and fan nozzles spray a straight line. Only misting nozzles and full cone nozzles are suitable in terms of their spray area because both spray over a full circular region, thereby maximizing the heat transfer area. Misting cannot be used because the pumping requirements are extremely high, which would cause restrictions in pump selection and add safety hazards. Therefore, a full cone spray nozzle was the best nozzle for this application.



Figure 3–6: UniJet Full Cone Spray Nozzle exploded view reproduced from Spraying Systems Co. Industrial Hydraulic Spray Products, Catalog 75 Metric Units (2015).

The nozzle that was selected is the stainless steel UniJet Full Cone Nozzle (1/4TT-SS+TG-SS0.3) from Spraying Systems Company. It has one of the lowest rated minimum flow rates available (160 mL/min at 1.5 bar). Its maximum rated flow rate is 390 mL/min at 10 bar. It has a spray angle of 50° at the minimum flow rate, but the angle increases to about 60° at higher flow rates. This is acceptable because only higher CPU powers, and therefore higher flow rates, require larger spray coverage. The nozzle also benefits from a replaceable spray tip and the ability to interchange different types of tips.

### 3.3.3 Spray Chamber

This section discusses the design process of the prototype spray chamber. First, a suitable material for the spray chamber is selected. Subsequently, the detailed designs of the spray chamber wall and then the spray chamber lid are given.

#### Material Selection

The material of the prototype spray chamber must be:

- transparent;
- corrosion and chemically resistant;
- rigid and mechanically strong;
- relatively inexpensive;
- able to be machined or 3D printed to a high degree of accuracy;
- and able to withstand temperatures much higher than 66.8 °C.

The requirement of transparency is important only for the prototype so that the spray can be analysed qualitatively during testing. At a minimum, the material must not melt or soften below a temperature of 66.8 °C because that is the predicted maximum temperature of the IHS surface. However, it should ideally be able to withstand temperatures much higher than 66.8 °C because a larger range of tested surface temperatures are desired during the simulated CPU tests.

The requirements of transparency and chemical resistance indicate that a plastic material is the best option. Two of the most common transparent plastic materials for prototyping and manufacturing are acrylic and polycarbonate, and both plastics meet all requirements. However, polycarbonate has a higher glass transition temperature (145 °C compared to 105 °C), and acrylic is more prone to chipping during manufacturing and during use; therefore, polycarbonate was selected for the prototype.

## Wall

The design of the spray chamber wall must facilitate certain functions. It must:

- allow for sufficient spray area on the surface of the IHS;
- create a hermetic seal with the surface of the IHS;
- connect to the LGA2011-3 ILM heat sink studs;
- create a hermetic seal with the spray chamber lid; and
- connect to the spray chamber lid.

Figure 3-7 shows the computer-aided design (CAD) of the prototype spray chamber wall. The first major design consideration is the size of the inner spray area of the wall. The maximum 5960X IHS surface dimensions are 36.74 mm × 40.5 mm. Although the IHS surface is slightly rectangular, a square design has been chosen to simplify all other aspects of the design and to allow the spray chamber to be placed in four orientations rather than just two. Also, a rectangular design has no advantage over a square design because the spray area is a circle. A square gasket has been chosen as the means of providing the hermetic seal between the wall and the IHS as well as between the wall and the lid. Ideally, a square O-ring would be used, but custom O-rings are expensive compared to hand-cut gaskets. The nominal height and width dimensions are both set to 3 mm. To avoid the possibility of the gasket slipping off any edge of the IHS, its outer dimensions have been chosen to be 35 mm × 35 mm. Taking into account the width of the gasket, this leaves a maximum spray area of 29 mm × 29 mm. Thus, the inner section of the wall design has these dimensions with rounded corners to facilitate manufacturing. Furthermore, the bottom and top of the wall

each feature four small extrusions at the corners of the inner perimeter. These extrusions are present to hold the gaskets in place.

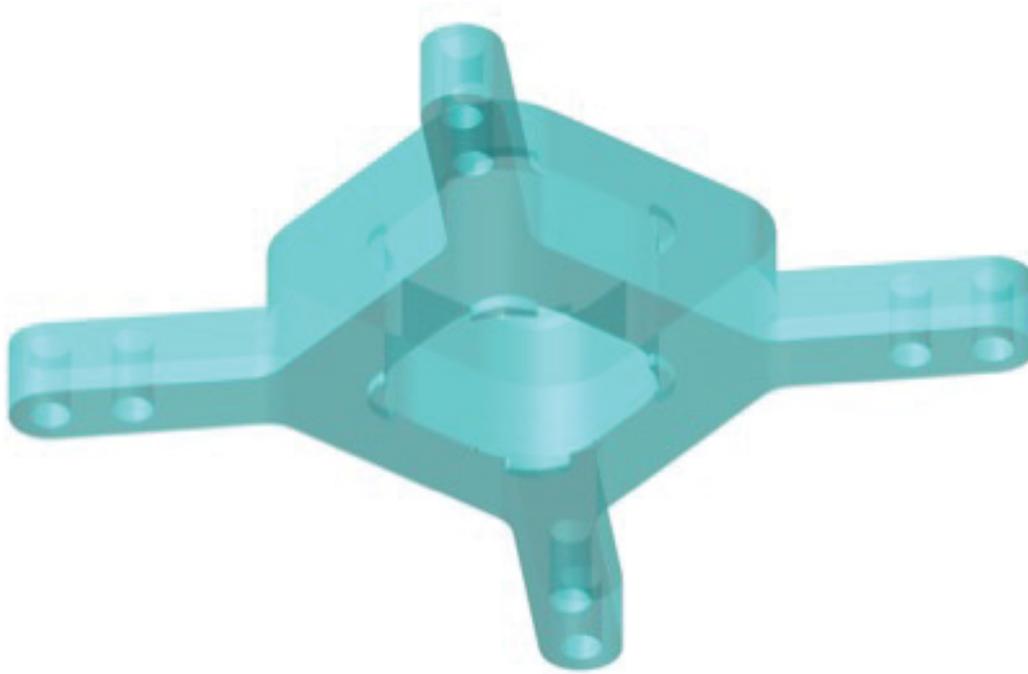


Figure 3-7: Spray chamber wall - CAD - bottom isometric view.

The wall design also features four arms that are extruding from its corners. The main purpose of these arms is to extend to the four ILM heat sink studs. The ILM heat sink studs accept M4 screws. The holes at the ends of the four arms are 5 mm in diameter, which allows a standard M4 shoulder bolt to fit through the holes and connect the wall to the ILM. Furthermore, these arms feature four inner holes that are used for mechanical fastening to the spray chamber lid.

All wall dimensions are fairly thick (nominally 10 mm) so that the polycarbonate cannot flex significantly and is less likely to break. The smallest critical thickness dimension is 2.5 mm, which is around the connection holes. The thickness of the gasket placer extrusions is smaller, but these extrusions only serve the purpose of ensuring proper placement of the gasket and do not support any significant load.

For the sake of spray observation, it is preferable that the height of the spray chamber lid be set such that the entire spray is observable. Given that the maximum possible spray diameter is 29 mm, and the maximum spray angle of the UniJet spray nozzle is 60°, the entire height of the of the spray chamber wall assembly, including the two gaskets, should nominally be set to 25.11 mm. The height has been rounded up to 26 mm for simplicity and to account for compression of the gaskets.

### **Lid**

The design of the spray chamber lid must:

- house and connect the spray nozzle;
- put the tip of the spray nozzle flush with the bottom of the lid;
- provide an outlet for the coolant;
- create a hermetic seal with the spray chamber wall; and
- connect to the spray chamber wall.

Figure 3–8 is the CAD representation of the prototype spray chamber lid design. The top of the spray chamber lid has a hole that is sized to install a through-wall fitting. The spray nozzle is connected to the fitting from inside the lid and inlet tubing connects to it from the outside. The fitting also creates a hermetic seal with the chamber. Another hole was bored in the side of the lid to support an additional through-wall fitting that serves as the coolant outlet. Moreover, the lid features four arms to facilitate connection with the chamber wall. The arms have holes at the ends that line up with the inner connection holes of the wall arms to allow a mechanical connection of the wall and lid with a fastener. Lastly, the height of the lid has been set so that the tip of the spray nozzle is flush with the bottom of the lid.



Figure 3–8: Spray chamber lid - CAD - top isometric view.

### 3.4 Assembly

Figures 3–9 and 3–10 show the assembled spray chamber. The spray chamber wall was machined by Xometry using computer numerical control (CNC). The spray chamber lid was 3D printed by Xometry using fused deposition modeling (FDM) rather than machining, and it was coated with Loctite EA E-20HP Epoxy to hermetically seal the surfaces. CNC machining is more expensive, but it is necessary for the chamber wall in order to maintain transparency in the spray section. Transparency is not absolutely necessary for the lid so the less expensive and faster FDM process could be used.



Figure 3-9: Spray chamber assembled - top isometric view.



Figure 3-10: Spray chamber assembled - bottom view.

## **CHAPTER 4**

### **Methods**

The following chapter describes the procedures undertaken to test the CPU cooling device using the experimental apparatus. It is split into separate simulated and real CPU sections to emphasize the procedural differences between the two experimental methods. Although the methods are recounted based on the experiments with the spray cooling device, the methods for new devices would be nearly identical.

#### **4.1 Simulated CPU Experiments**

This section describes the details of the experiments performed using the prototype CPU cooling device attached to the simulated CPU. Each series of experiments began with a setup of the apparatus and its control and data acquisition system. Then, the experiments were conducted at each flow rate and simulated CPU power setting of interest. At the end of a series of experiments, the system was safely shutdown, and the data files were obtained from the RPi.

##### **4.1.1 Setup**

Before starting any setup procedures, it was important to ensure that the main power bar connected to the wall outlet was powered down for the safety of the user and electrical system. After the power was turned off, the main reservoir was inspected to check that there was enough coolant. The liquid level of the coolant had to be within a few centimeters of the top of the reservoir to ensure that the pump never ran dry.

With the coolant in order, the CPU cooling device was connected to the simulated CPU and the fluid loop. In the case of the spray cooling device, the square gasket was placed on the bottom of the spray chamber wall, and the spray chamber was then set on the copper block surface of the simulated CPU such that the heat sink screw holes of the cooling device aligned

with those of the simulated CPU case. For other types of cooling devices, a TIM may need to be applied between the bottom of the cooling device and the copper block, but this is not necessary for the spray cooling design. Next, the cooling device was mechanically fastened to the simulated CPU. For the spray cooling device, the seal between the spray chamber and the copper block was very important because improper seals led to leaks. After making a new connection, the seal was monitored closely during the first few experiments to ensure that there were no leaks. Once the cooling device was properly connected to the simulated CPU, the mating halves of the fluid loop's quick-disconnect fittings were connected to the inlet and outlet ports of the device. Then, the device and the simulated CPU case were inverted and placed in the front stand of the experimental apparatus platform. Finally, the fluid loop tubing lines were connected using the quick-disconnect fittings.

After the fluid loop preparation was complete, the electronics and circuits of the control and data acquisition system were inspected for any potential issues, such as loose connections. Before launching the application, the pump was set to analog mode so that the RPi could control it. After launching the application and the graphical user interface appeared, the system parameters were set and verified. If the remainder of the startup process of the program did not produce any error messages in the log panel, then the system was ready to start experiments.

#### **4.1.2 Experiment**

The procedure for performing the experiments was straightforward because everything could be controlled directly from the user interface and the data acquisition was automatically handled by the application. Each testing point of the experiment had two control parameters: the simulated CPU power and the coolant flow rate. For each power setting, a full series of flow rates were tested starting with the highest flow rate. The flow rate was limited to prevent the pump outlet pressure from exceeding 4 bar because there were occasionally leakage problems with tubing line connectors at pressures higher than 4 bar. The lowest flow rate tested was usually the minimum flow rate required to achieve a proper spray

pattern (observed through the transparent spray chamber walls), but it had to be set higher at higher powers so that the simulated CPU did not overheat.

Starting the experiment in the application turned on the pump, the fans and the breaker of the simulated CPU; however, the power to the simulated CPU was always initially set to 0 W. Next, the pump flow rate was changed to the desired setting. The spray pattern was then observed qualitatively through the walls of the spray chamber to ensure that it was properly impinging on the simulated CPU surface. At this point, the power to the simulated CPU could be set to the desired value. At some point after changing the power setting, the alarm would beep once to indicate that the simulated CPU was in a steady state on a time scale of 2 minutes (i.e. simulated CPU temperatures were not 1 °C different compared to values 2 minutes ago) and then twice to indicate a steady state on a time scale of 4 minutes. It took between 5 and 25 minutes to reach steady state between testing points. Once both steady state indicators were true, the measurements could begin. To indicate a measurement point in the data file, the test identifier was set in the application to a unique integer identifier for that measurement point. After a measurement was complete, which would nominally be after 4 minutes, the test identifier was set back to the default, 0. The next testing point's flow rate and/or power were then set.

#### **4.1.3 Shutdown and Obtaining Data**

The shutdown process for the experiments was important because serious damage to the CPU cooling device was possible if the shutdown was not handled carefully. The simulated CPU heating element, insulation and case constitute a large thermal mass, and a significant amount of internal energy remains in the system after the power to the heating element is turned off. Therefore, once a series of experiments was concluded, the heat was turned off, and the coolant flow rate was increased to its maximum value (without raising the pump outlet pressure above 4 bar). Then, the temperature of the simulated CPU was monitored until it was deemed to be low enough to stop the spray cooling. The simulated CPU was considered safe when the surface temperature was below 27 °C, and the case temperature

was below 57 °C. This process could take over 20 minutes. Once safe, the flow rate was set to 0 mL/min, and the experiment was ended using the application controls. When the user quit the application, the remaining connections would automatically shutdown safely, and the data, parameter and log files would be finalized.

## 4.2 Real CPU Experiments

Although there were many similarities in the real CPU experimental methods compared to the simulated CPU methods, there were also important differences, and those differences are highlighted in this section. Like the previous section, this section is subdivided into the setup, experimental and shutdown procedures.

### 4.2.1 Setup

The setup procedure for the real CPU experiments was very similar to that of the simulated CPU experiments. The major difference was in the connection of the cooling device. In the real CPU case, the cooling device had to be connected to the CPU and CPU socket of the HF320 server rather than the simulated CPU. To accomplish this, the server stand was first set to the appropriate height for the cooling device being tested. For the spray cooling device, the bottom of the stand's platform had to be at least 15 cm above the top of the radiator. Then, the top cover of the server's case was removed, and the server was placed and secured in the stand (in an inverted manner for the spray cooling device). Next, the assembled cooling device was connected to the CPU socket and sealed to the CPU's IHS using the four M4 shoulder screws and the square gasket, respectively. The remainder of the setup procedure was nearly identical to that of the simulated CPU, except further connections had to be made to server. An additional monitor, keyboard and mouse set was connected to the server, and the two power cords of the server were connected to a power bar that was plugged into the modified KAW and controllable breaker outlet (in place of the simulated CPU plug).

### 4.2.2 Experiment

In contrast to the testing points of the experiments with the simulated CPU, each testing point of the real CPU experiments had only one parameter of control: the CPU frequency (clock speed). This parameter is akin to the simulated CPU power. The coolant flow rate, which was the second parameter of control for the simulated CPU experiments, was instead fixed to the maximum rate tested in those experiments because the effect of changing flow rate was already investigated and high-performance computers would, realistically, only use the highest flow rate. Thus, when the pump was turned on at the beginning of the experiment, the flow rate was immediately changed to its maximum value and was not modified again.

Table 4-1: UEFI overclocking settings.

Cache Multiplier	38-39X
Cache Voltage	1.3 V
CPU Input Voltage	1.92 V
RAM Frequency	2666 MHz
CPU Load-line Calibration	LVL 1
Fixed CPU VRM Switching	500 kHz
Power Phase Control	Extreme
Power Duty Control	Extreme
Power Current Capability	110%
Power Thermal Control	130
Integrated VR Fault Management	Disabled
Integrated VR Efficiency Management	High Performance
Turbo Mode	Enabled
Enhanced Intel CPU Speedstep Technology	Disabled
Intel Adaptive Thermal Control	Enabled
Max Allowed Core Temperature	Auto
Cores	ALL 8
Hyper-Threading	Disabled

With the pump running at the appropriate flow rate, the server was then turned on by switching on the server breaker through the application and pressing the physical button of the server. At this point, the UEFI was launched during boot. When the UEFI graphical user interface launched, the overclocking settings were modified by navigating the available menus with the keyboard. To enable overclocking, it was first ensured that certain options

were properly set (see table 4-1). Second, the CPU frequency was set by modifying the CPU core ratio and raising the CPU core voltage to its corresponding minimum value (see table 4-2). Once the UEFI settings were ready, the changes were saved, and the server was restarted. At this time, the operating system was allowed to start. Next, the OCCT 4.2.0 and HWiNFO64 programs were opened. The OCCT test type ('CPU:OCCT') was then selected and the options were set to 'Automatic,' duration of 18 minutes, idle periods of 1 minute at the beginning and end, and 'Large Data Set' test mode. When the test was started with the OCCT 'ON' button, the data recording button of HWiNFO64 was pressed to save the server sensor data to file, and a separate timer was set for 12 minutes. After those first 12 minutes, the CPU was considered to be in steady state, so the test identifier was set to its unique integer identifier for that measurement point using the RPi application. After a measurement period of 4 minutes, the test identifier was set back to 0. Finally, the OCCT test was allowed to complete, and the HWiNFO64 data recording was stopped. If the test finished without error, the cooling system was considered to have passed at that CPU frequency. However, if the server crashed or produced an error at any point of the test, the cooling system was considered to have failed at that frequency. When the test was complete, the server was restarted, and the next CPU frequency was set in the UEFI.

Table 4–2: UEFI settings for specified CPU frequency.

CPU Frequency [GHz]	Core Ratio [X]	Core Voltage [V]
3.0	30	1.000
3.1	31	1.000
3.2	32	1.000
3.3	33	1.000
3.4	34	1.000
3.5	35	1.015
3.6	36	1.030
3.7	37	1.045
3.8	38	1.060
3.9	39	1.075
4.0	40	1.090
4.1	41	1.112
4.2	42	1.121
4.3	43	1.150
4.4	44	1.170
4.5	45	1.214
4.6	46	1.252
4.7	47	1.315

### 4.2.3 Shutdown and Obtaining Data

Prior to shutting down the server during the last experiment, the server sensor data files were retrieved via a USB drive. Unlike the simulated CPU, the real CPU did not pose a threat to the cooling device after shutdown because the real CPU constitutes a significantly smaller thermal mass and does not reach temperatures above 100 °C anywhere in the package. Thus, shortly after turning off the server, the thermal state was considered safe, so the RPi program was closed, and the RPi files were collected following the same procedure as the simulated CPU experiments.

## CHAPTER 5

### Results

The results of the experiments using the spray cooling device prototype with the two tested coolants are discussed in the present chapter. These initial results determine the effectiveness of the prototype as a processor cooling device, and also validate the experimental apparatus as a suitable platform for testing processor cooling devices. The first section is concerned with the fluid flow characteristics using the cooling device over a series of flow rates. The next section presents the results of the simulated CPU experiments with the spray cooling device, and the final section describes and benchmarks the results of the real CPU experiments. An uncertainty analysis for reported measurements is presented in Appendix B.

#### 5.1 Fluid Flow Rate and Pressure Experiments

After attaching the simulated CPU to the spray cooling device, but prior to the simulated CPU experiments, a series of fluid loop tests were conducted. The purpose of these tests was to operate the pump over a range of flow rates to monitor the pump outlet pressure and check for any leaks. These tests were performed with distilled water as the coolant, using the same methods as the simulated CPU experiments, except that no power to the heating element was applied.

The initial tests helped to find and eliminate leaky connections. Eventually, the system was capable of running at pump outlet pressures up to 4 bar without a leak. Finally, a full series of flow rates from 0 mL/min to 255 mL/min were tested in sequence. The results of these tests, in terms of the pressure drop across the spray nozzle and the flow rate, are plotted in figure 5–1. This data was then fit, using least squares regression, to a simple

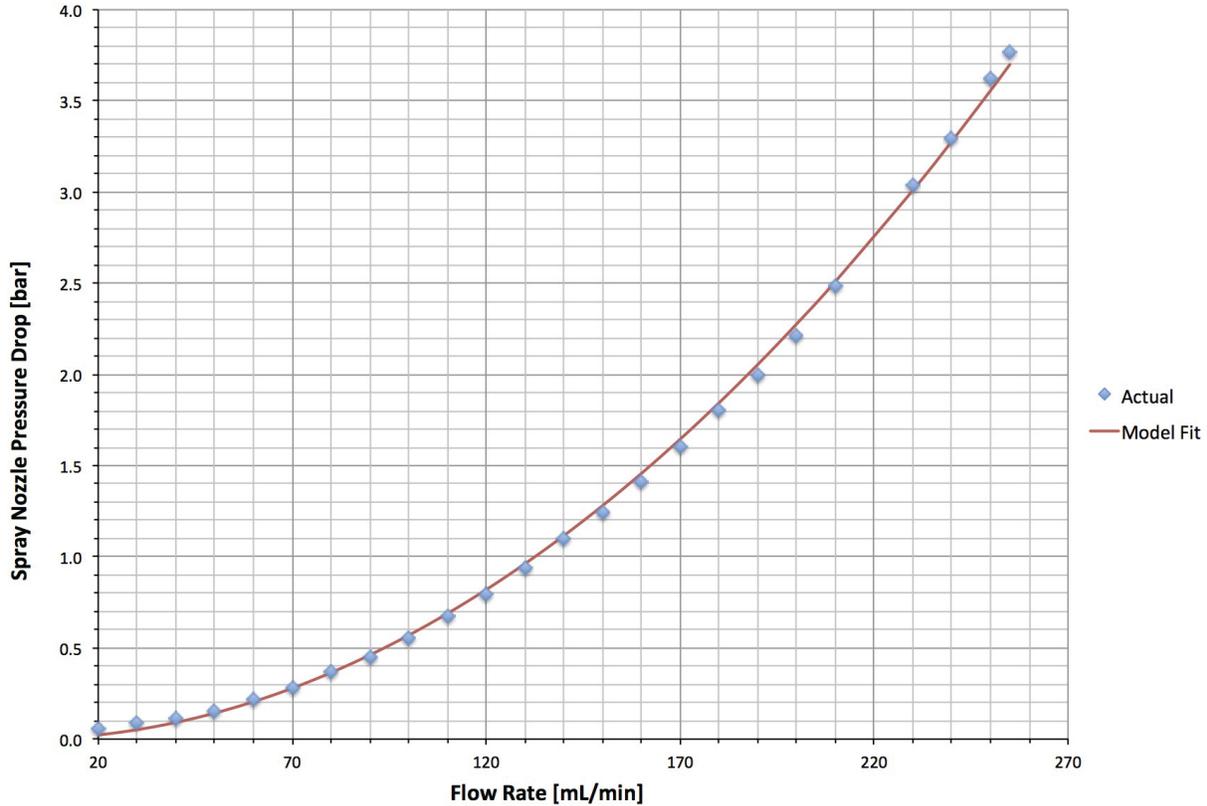


Figure 5–1: Spray nozzle pressure drop as a function of flow rate for distilled water.

model equation for spray nozzle pressure as a function of flow rate:

$$\Delta p_{nozzle} = \left( 5.69 \times 10^{-5} \frac{\text{bar}}{(\text{mL}/\text{min})^2} \right) \left( \frac{\rho_l}{\rho_{water}} \right) Q^2, \quad (5.1)$$

where  $\Delta p_{nozzle}$  is the pressure drop across the spray nozzle,  $5.69 \times 10^{-5} \text{ bar}/(\text{mL}/\text{min})^2$  is an experimentally determined constant,  $\rho_l$  is the coolant density,  $\rho_{water}$  is the density of liquid distilled water, and  $Q$  is the coolant flow rate. This model function is plotted alongside the data in figure 5–1. The data and function closely match the spray nozzle manufacturer’s reported values. Furthermore, this function is used in the control and data acquisition application to make expected pump outlet pressure calculations. A safety case (from the safety monitor described in section 2.4.5) is in place to compare the expected pressure to the actual pressure, and the program outputs a warning if they differ significantly.

## 5.2 Simulated CPU Experiments

The simulated CPU experiments were the first heated tests conducted using the experimental apparatus. This section describes the results of those experiments, and is divided into two subsections, the first pertaining to the single-phase coolant tests, and the second pertaining to the two-phase coolant tests.

### 5.2.1 Single-Phase Coolant: Distilled Water

The first set of heated tests were conducted with distilled water and the simulated CPU oriented in an inverted manner. The pressure inside the spray chamber was equal to atmospheric pressure (as confirmed by the pressure sensor at the spray chamber exit), and the ambient temperature was maintained at an average of  $21.5 \pm 0.2$  °C. The temperature of the coolant entering the spray chamber and nozzle was approximately equal to the ambient temperature, so the subcooling (temperature below the saturation temperature) was about 78 °C.

The surface temperature of the CPU's IHS and the heat transfer rate across that surface are the two variables of most interest in processor cooling applications. In the case of the simulated CPU experiments, the surface temperature is estimated using the RTD sensor embedded in the copper block of the simulated CPU's heated surface. Although the power to the heating element is recorded, that power is not an accurate estimation of the heat transfer rate across the surface because it does not account for losses. Therefore, the heat transfer rate was instead determined using conservation of energy based on the sensible heat transfer, and assuming steady state, as well as no work, and constant potential and kinetic energy:

$$q = \rho_l c_{p,l} Q (T_{out} - T_{in}), \quad (5.2)$$

where  $q$  is the heat transfer rate across the surface,  $T_{out}$  is the spray chamber outlet temperature, and  $T_{in}$  is the spray chamber inlet temperature. This equation does not account for the heat loss through the spray chamber walls, but it serves as a good minimum estimate of the

heat transfer rate. Moreover, it is not suitable for two-phase fluid flow. Therefore, a model<sup>1</sup>, based on the simulated CPU power and temperatures, was developed to estimate the heat transfer rate using the distilled water experimental data in the FC-3284 experiments.

---

<sup>1</sup> See Appendix A for the derivation of the heat transfer model equation

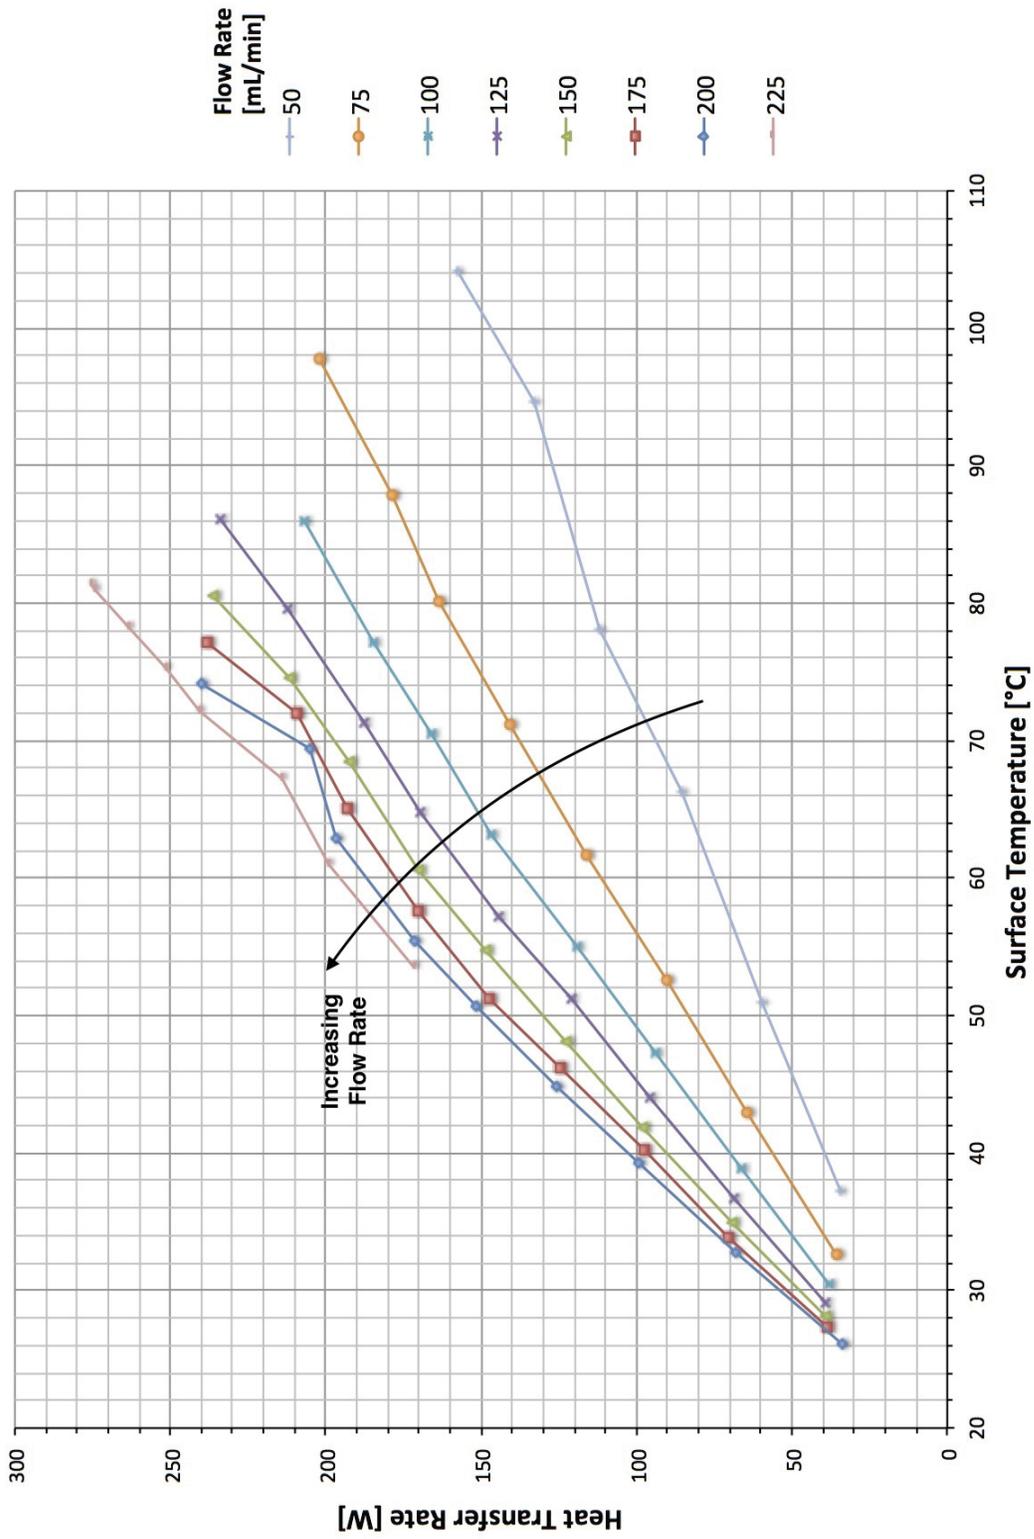


Figure 5-2: Heat transfer rate as a function of surface temperature and flow rate for the simulated CPU using distilled water as the coolant.

The heat transfer rates for the distilled water experiments are plotted as a function of surface temperature for each coolant flow rate in figure 5–2. (Although the heat transfer rate is the independent variable in these experiments, it is plotted on the vertical axis with the surface temperature on the horizontal axis, because this is the standard convention in the field.) The general trend of the data is a linear increase in heat transfer rate with increased surface temperature. Moreover, as one would expect, the heat transfer increases at higher flow rates. The best performing and highest flow rate tested was 225 mL/min, and the highest heat transfer rate achieved with this flow rate was 275 W at a surface temperature of 81.3 °C. However, the heat transfer rate at the maximum recommended temperature of the CPU’s IHS surface (66.8 °C) is of more importance to CPU cooling. For this case, the heat transfer rate is approximately 215 W at a flow rate of 225 mL/min. Furthermore, it is clear that increasing the flow rate above 225 mL/min would improve the heat transfer rate at that critical temperature; however, increases in flow rate at that magnitude have a smaller impact on heat transfer compared to the increases in flow rate at the lower flow rate magnitudes (e.g. increasing from 50 mL/min to 75 mL/min).

### **5.2.2 Two-Phase Coolant: FC-3284**

Following the experiments with the distilled water, a set of experiments were conducted with the FC-3284 engineered coolant. Once again, the simulated CPU was oriented in an inverted manner, and the pressure inside the spray chamber was effectively equal to atmospheric pressure. Moreover, the ambient temperature was fairly constant ( $21.1 \pm 0.3$  °C), which resulted in subcooling of approximately 29 °C. The execution of these experiments was similar to that of the distilled water experiments; however, the results required a more in-depth analysis due to the two-phase nature of the cooling.

A plot of the heat transfer rate as a function of surface temperature and flow rate is presented in figure 5–3. As mentioned in the previous section, the sensible heat transfer alone is not an accurate representation of the total heat transfer rate in this case due to the additional two-phase cooling. Therefore, a model was developed to estimate the heat

transfer rate based on the input data of the simulated CPU power and temperatures (see Appendix A). The final form of that model is:

$$q = \frac{q_{total} [1.995 \frac{K}{W} + (4.112 \frac{K^{\frac{5}{4}}}{W}) (T_{case} - T_{amb})^{-\frac{1}{4}}] - (T_{surf} - T_{amb})}{3.368 \frac{K}{W} + (4.112 \frac{K^{\frac{5}{4}}}{W}) (T_{case} - T_{amb})^{-\frac{1}{4}}}, \quad (5.3)$$

where  $q_{total}$  is the total heat generated by the simulated CPU, which is equal to the average input power of the simulated CPU ( $P_{ave}$ ), and where  $T_{amb}$  is the ambient temperature.

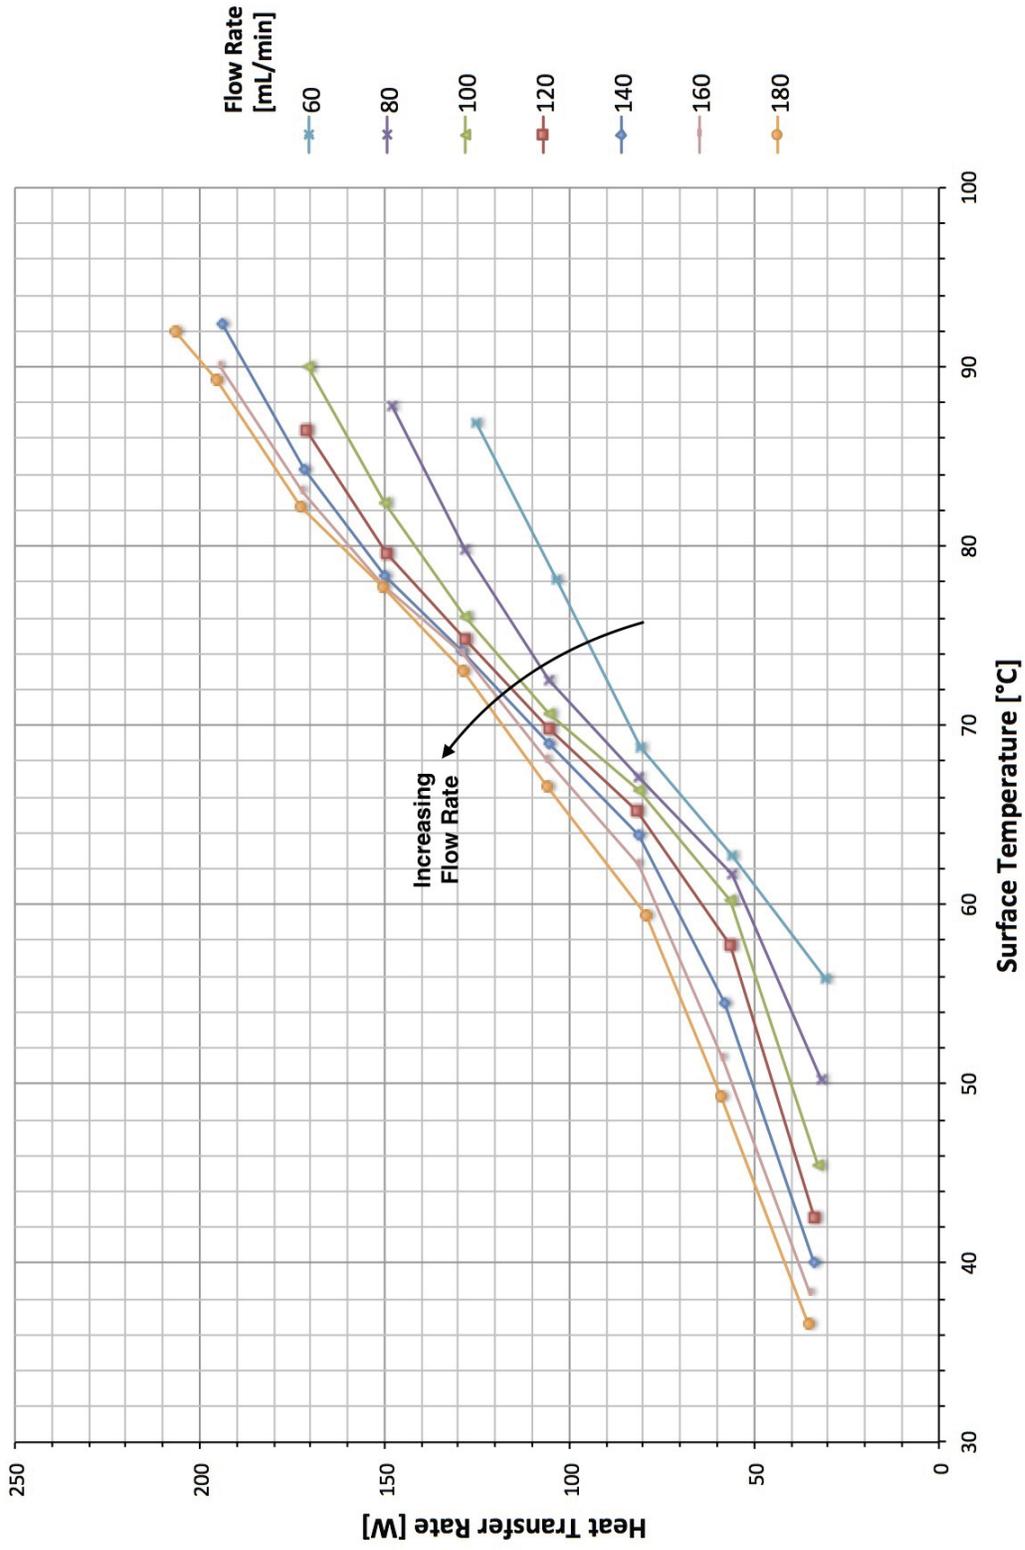


Figure 5-3: Heat transfer rate as a function of surface temperature and flow rate for the simulated CPU using FC-3284 as the coolant.

The trend in figure 5–3 of increasing heat transfer rate with increasing surface temperature is much less linear when compared to the single-phase coolant (figure 5–2). Below the boiling point of the coolant ( $T_{sat}=50\text{ }^{\circ}\text{C}$ ), in the single-phase regime, where only sensible heat transfer occurs, the trend appears to be linear (there are not enough data points to confirm this). However, in the two-phase regime after the boiling point, there is an inflection point and the heat transfer rate increases more rapidly with surface temperature. This is due to the additional latent heat transfer associated with the vaporization of the coolant. Further on in the two-phase regime, there is a second inflection point that indicates a transition towards the critical heat flux, where the coolant is no longer capable of increased heat transfer. However, the CHF was not reached for any of the flow rates, so the exact temperature and heat transfer rate in which that occurs for each flow rate was not determined.

Using FC-3284 as the coolant results in lower heat transfer rates than using distilled water. The highest heat transfer rate was 206 W at a surface temperature of 91.9 °C and flow rate of 180 mL/min. Furthermore, at the IHS’s maximum allowable surface temperature (66.8 °C), the maximum heat transfer was only about 106 W. There are two main reasons that the cooling is not as effective with the FC-3284. Firstly, the specific heat capacity of FC-3284 is one quarter of the specific heat capacity of water; thus, each gram of coolant is absorbing much less sensible heat per degree Celsius. Secondly, latent heat transfer is not occurring until surface temperatures relatively close to the maximum allowable temperature of the IHS.

Figures 5–4, 5–5 and 5–6 give further insight into the two-phase heat transfer characteristics observed in the FC-3284 experiments. Figure 5–4 plots the sensible heat transfer rate as a function of surface temperature for different coolant flow rates. The sensible heat transfer rate was calculated using the same equation that was used for distilled water (equation 5.2). The heat transfer curves of this plot follow the same trend in the single-phase regime as the total heat transfer curves, as would be expected. However, the curves reach a heat transfer rate plateau in the two-phase regime. The onset of this plateau represents the point at

which all of the coolant has reached its saturation temperature. At this point, sensible heat transfer in the liquid coolant has reached its maximum, and any additional heat transfer is gained due to latent heat transfer. Further sensible heat transfer in the coolant vapour is not observed in this plot, which was confirmed by the fact that the coolant temperature at the outlet of the spray chamber never exceeded the saturation temperature.

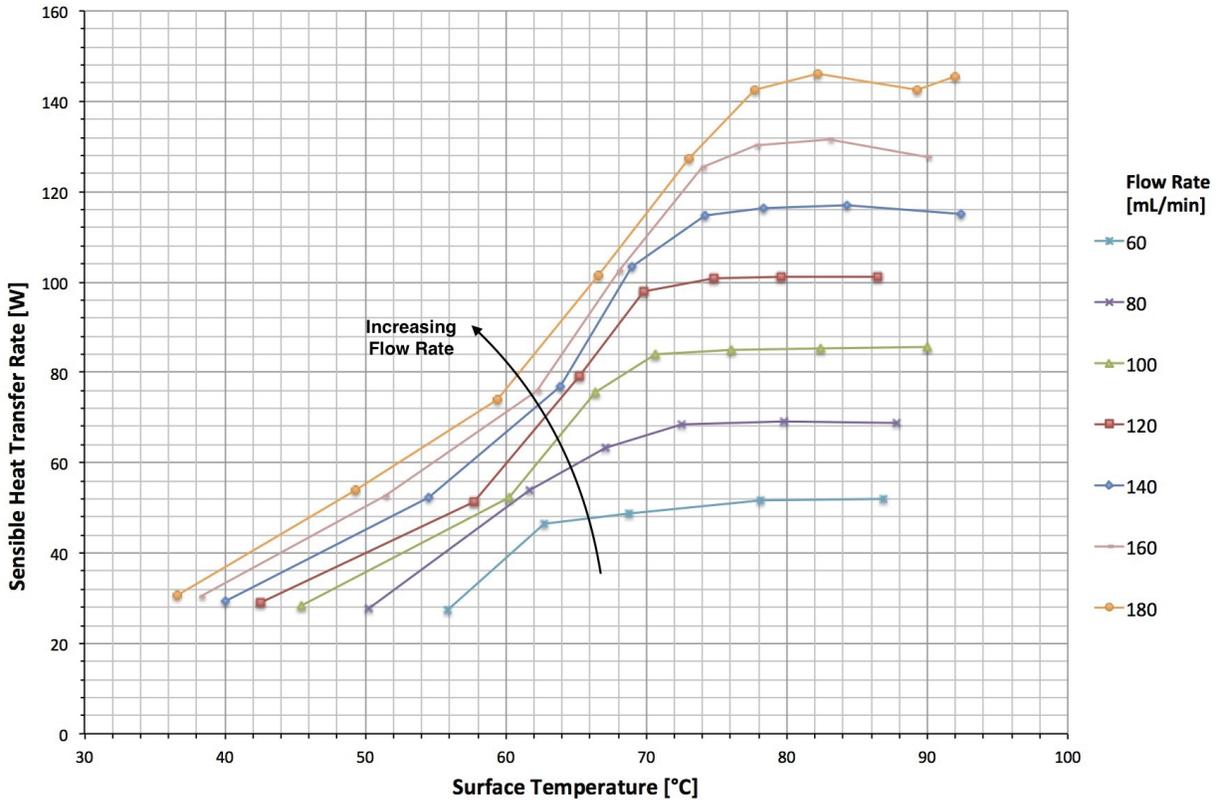


Figure 5–4: Sensible heat transfer rate as a function of surface temperature and flow rate for the simulated CPU using FC-3284 as the coolant.

The vapour quality ( $\chi$ ), which is the mass fraction of coolant that is vapour, was calculated using:

$$\chi = \frac{(q - q_{sens})/\lambda}{\rho_l Q}, \quad (5.4)$$

where  $q_{sens}$  is the sensible heat transfer rate (calculated using equation 5.2) and  $\lambda$  is the latent heat of vaporization. Figure 5–5 plots the vapour quality as a function of surface temperature and flow rate. The curves follow a standard trend: there is negligible vaporization until

a flow-rate-dependent surface temperature above the boiling point, after which the mass fraction of coolant being vaporized starts increasing relatively linearly with increased surface temperature. Moreover, more coolant is vaporized at the lower flow rates compared to the higher flow rates at a given surface temperature. One noteworthy observation is the fact that there is negligible vaporization of the coolant, even at the lowest flow rate, until surface temperatures over 10 °C above the boiling point. While a superheated surface (surface with a temperature above the coolant saturation temperature) is typically required to initiate vaporization, a superheat over 10 °C is high. At the highest flow rate, vaporization only starts occurring at temperatures above 78 °C, which is 28 °C above the boiling point. Even at the lowest flow rate, the highest fraction of coolant vaporized was only approximately 0.4, which signals poor spray efficiency.

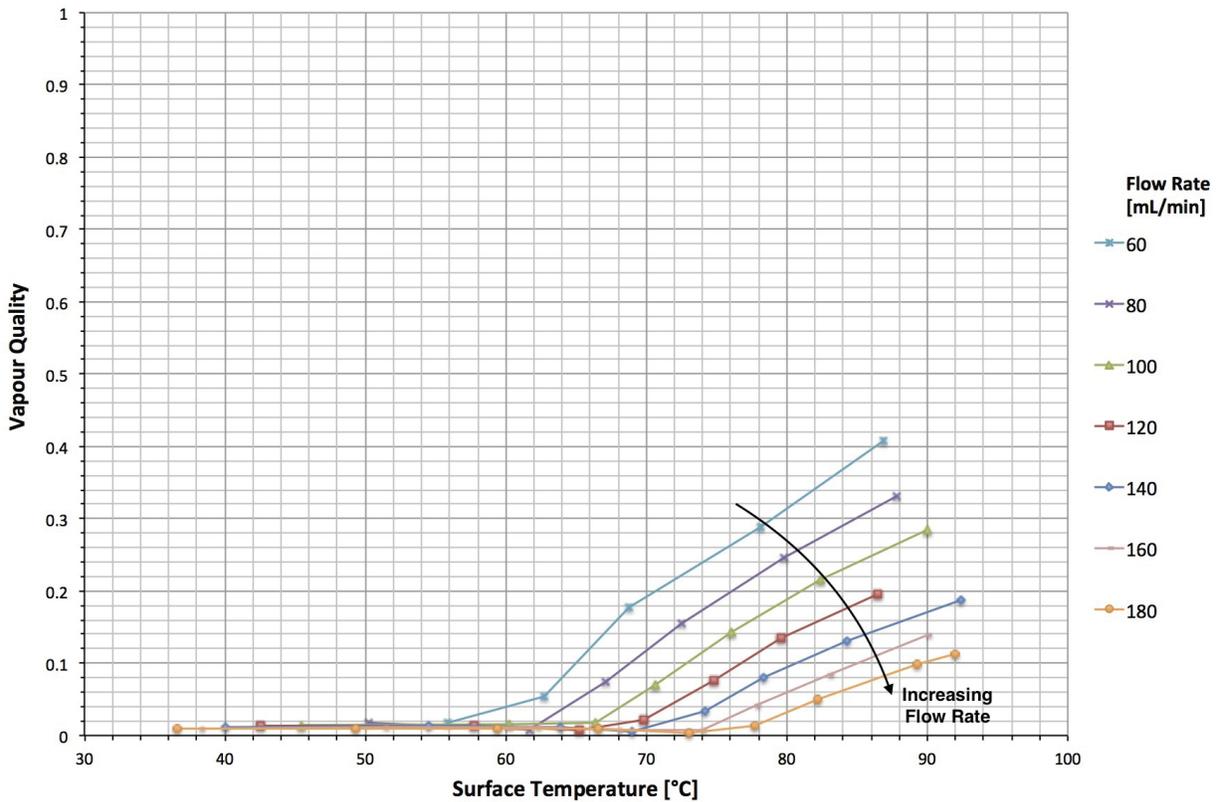


Figure 5-5: Vapour quality as a function of surface temperature and flow rate for the simulated CPU using FC-3284 as the coolant.

Spray efficiency ( $\eta$ ) is an indicator of the utilization of the total “energy-absorbing ability” of the coolant, and is calculated as follows:

$$\eta = \frac{q}{\rho_l Q (c_{p,l}(T_{sat} - T_{in}) + \lambda)}. \quad (5.5)$$

Figure 5–6 plots the spray efficiency as a function of surface temperature for the tested flow rates. A spray that is fully efficient ( $\eta=1$ ) vaporizes all of the coolant. Thus, the spray efficiency is strongly related to the vapour quality, which is apparent when comparing figures 5–5 and 5–6. All flow rates follow a single curve in the single-phase regime, but they gradually break away from that curve in the two-phase regime. Since the vaporization is greatest at the lower flow rates, the lower flow rates also have the highest efficiency for a given surface temperature. The maximum efficiency achieved was 0.54 at a flow rate of 60 mL/min and surface temperature of 86.9 °C. Below the maximum allowable temperature of the IHS, the efficiencies are below 0.31; thus, the FC-3284 is not used very efficiently in this application of processor cooling.

Based on these findings, spray cooling with FC-3284 is likely not suitable for processor cooling using the current design and flow rates. Further increasing the flow rate would improve overall heat transfer rates, but, at higher flow rates, latent heat transfer would be reduced, and the coolant would effectively be operating as a single-phase coolant; therefore, it would be preferable to use a superior single-phase coolant, such as water. A better option is to use a different engineered coolant with a lower boiling point, such as Novec 7000 from 3M ( $T_{sat}=34$  °C). Since this coolant has similar properties to FC-3284, its use would hypothetically have the same effect as shifting the FC-3284 heat transfer curves left by 16 °C, which would result in higher heat transfer rates falling into the required temperature range.

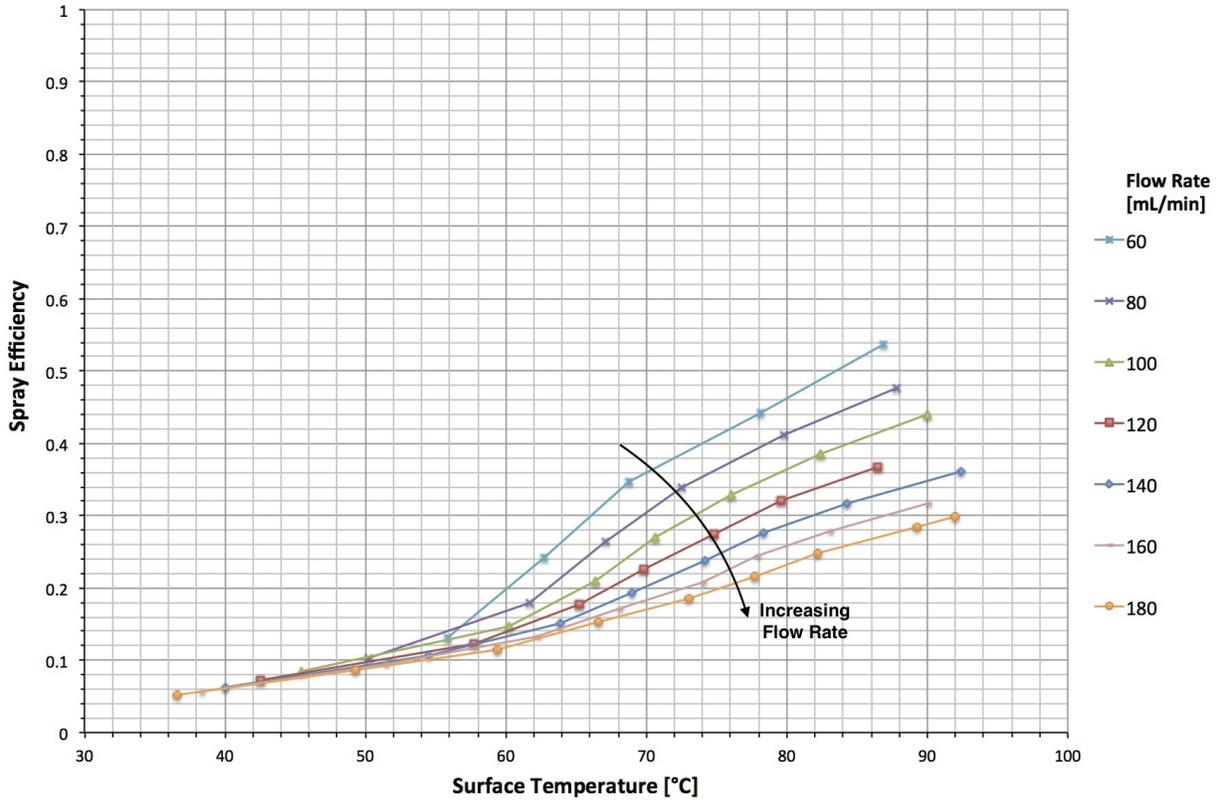


Figure 5-6: Spray efficiency as a function of surface temperature and flow rate for the simulated CPU using FC-3284 as the coolant.

### 5.3 Real CPU Experiments

After concluding the simulated CPU experiments, the experimental apparatus, spray cooling device and coolants were tested using the 5960X CPU in the HF320 server. However, prior to performing the distilled water and FC-3284 spray cooling experiments, a series of tests were undertaken using the (standard) Asetek microchannel cooling system installed in the HF320 server to benchmark the processor’s cooling. These tests were conducted following the same methods of section 4.2; but, the server’s top lid was not removed, and the server was not inverted in the stand. Furthermore, the fluid loop of the Asetek cooling system was used rather than that of the experimental apparatus because it consists of an all-in-one system that is not easily disassembled. Therefore, fluid loop sensor data recorded on the RPi could not be analyzed for these tests.

All experiments were conducted at similar ambient temperatures. The ambient temperatures for the Asetek, distilled water and FC-3284 tests were  $20.1\pm 1.0$  °C,  $20.9\pm 0.3$  °C and  $21.0\pm 0.1$  °C, respectively. Therefore, the subcooling for distilled water was approximately 79 °C, and it was approximately 29 °C for FC-3284. Moreover, as with the simulated CPU, the distilled water and FC-3284 experiments were conducted with the spray chamber (and server) inverted, and the pressure in the spray chamber remained approximately equal to atmospheric pressure.

The most important results for the real CPU experiments are obtained from knowing whether the cooling system passes or fails at each CPU frequency of interest. The criteria for passing are that the i) server successfully passes the OCCT stress test for its 18 minute duration without crashing or producing an error, and ii) CPU core temperatures do not rise above 100 °C (since that is approaching the maximum allowable core temperature of 105 °C). The results are summarized in table 5–1. In this table, ‘N/A’ indicates that the cooling system was not tested at that frequency. Originally, the lowest tested frequency was going to be 3.5 GHz; however, the system cooled with FC-3284 was hardly capable of passing at 3.5 GHz, so additional lower frequency tests were added for that coolant. There was no benefit to performing these lower frequency tests with the distilled water and Asetek system, as they would certainly have passed. As noted in table 5–1, the Asetek system performed best and was able to pass every test up to 4.7 GHz. The distilled water spray cooling device was capable of passing the stress tests up to 4.5 GHz. Therefore, the spray cooling design, at its current flow rate, results in lower processor cooling relative to the state-of-the-art Asetek microchannel system. However, it should be noted that the Asetek system operates at flow rates around 1,500 mL/min, while the spray cooling device only operates at a flow rate of 225 mL/min. Thus, by slightly modifying the configuration, the spray cooling device may result in improved cooling at higher flow rates.

To analyze the cooling system’s failures, it is important to begin by examining the cores of the CPU, which are the main sources of heat, and are thus the hottest components of

Table 5–1: Pass or fail results of the real CPU experiments for the Asetek, distilled water spray and FC-3284 spray cooling systems.

CPU Frequency [GHz]	Asetek	Distilled Water	FC-3284
3.0	N/A	N/A	Pass
3.1	N/A	N/A	Pass
3.2	N/A	N/A	Pass
3.3	N/A	N/A	Pass
3.4	N/A	N/A	Pass
3.5	Pass	Pass	Pass
3.6	Pass	Pass	Fail
3.7	Pass	Pass	Fail
3.8	Pass	Pass	Fail
3.9	Pass	Pass	Fail
4.0	Pass	Pass	Fail
4.1	Pass	Pass	Fail
4.2	Pass	Pass	Fail
4.3	Pass	Pass	Fail
4.4	Pass	Pass	Fail
4.5	Pass	Pass	Fail
4.6	Pass	Fail	Fail
4.7	Pass	Fail	Fail

the CPU package. Furthermore, the stability of the CPU is strongly dependent on the core temperatures, and can cause the CPU to produce an error or crash at higher core temperatures. The HWiNFO64 software records data from various sensors embedded into the computer system, including the temperatures for each of the eight cores of 5960X CPU. The average and maximum core temperatures as functions of CPU frequency for each of the three cooling systems are plotted in figures 5–7 and 5–8. The accuracy of this sensor data is not known; therefore, the data is best only used qualitatively.

The trends in the average core temperature and maximum core temperature plots are nearly identical, with the maximum temperature being an average of 6 °C above the average value. At the highest CPU frequencies, that difference is upwards of 8 °C, which indicates that the variations in core temperature become more volatile at higher frequencies. Another trend

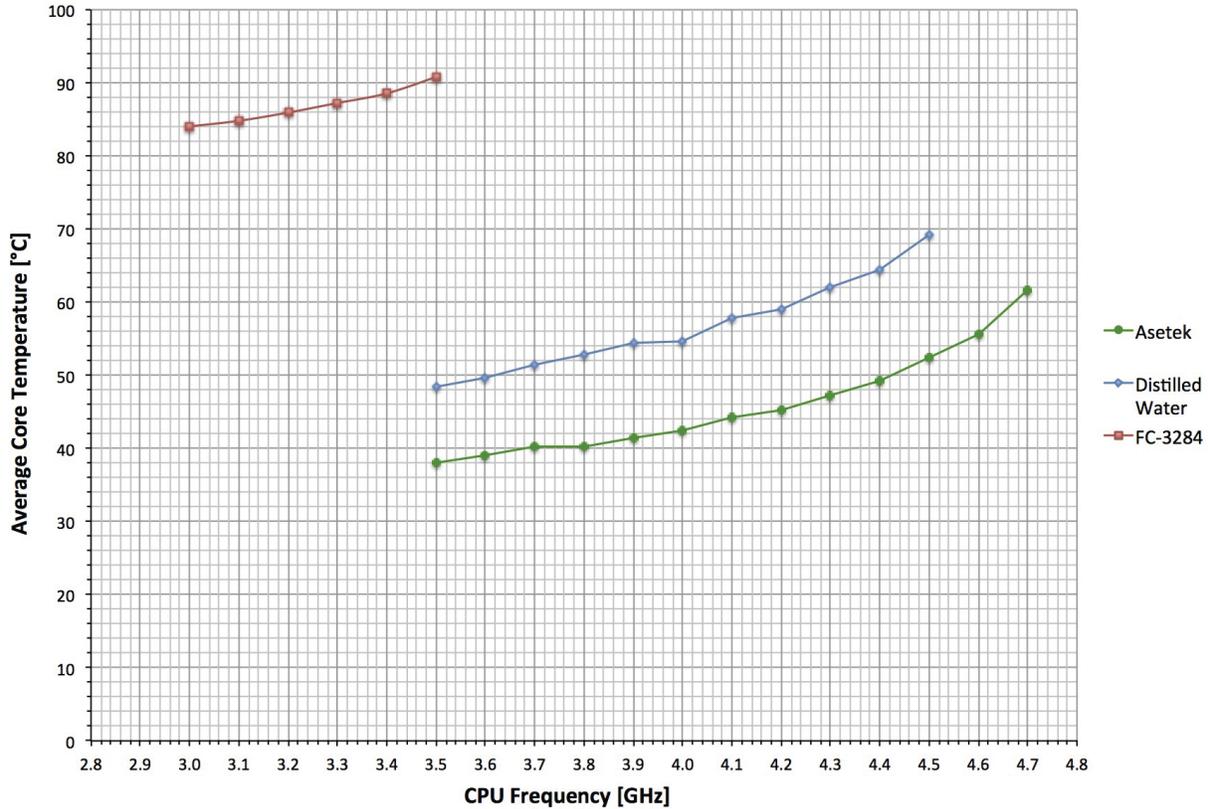


Figure 5–7: Average CPU core temperature comparison for the real CPU experiments using the Asetek, distilled water spray, and FC-3284 spray cooling systems.

apparent in all cases is that core temperatures rise exponentially with increased frequency. This trend is due to the increased CPU power required to operate at elevated frequencies.

As stated previously, elevated core temperatures can result in processor error or failure. However, this statement appears to be contradicted by the fact that the FC-3284 tests were capable of passing with core temperatures over 90 °C, while the distilled water tests at 4.6 GHz failed with core temperatures around 80 °C. The explanation is that the higher CPU power and frequencies encountered in the distilled water tests result in greater core temperature volatility, which is only partially sensed by the core temperature sensors. The sensor measurements are made on a macroscale, and the nanoscale local temperature variations in the cores are not measured; thus, the full degree of volatility is not captured by the temperature sensors, and the maximum temperatures at the transistor level may be excessive and causing errors. Therefore, the maximum average core temperature limit is

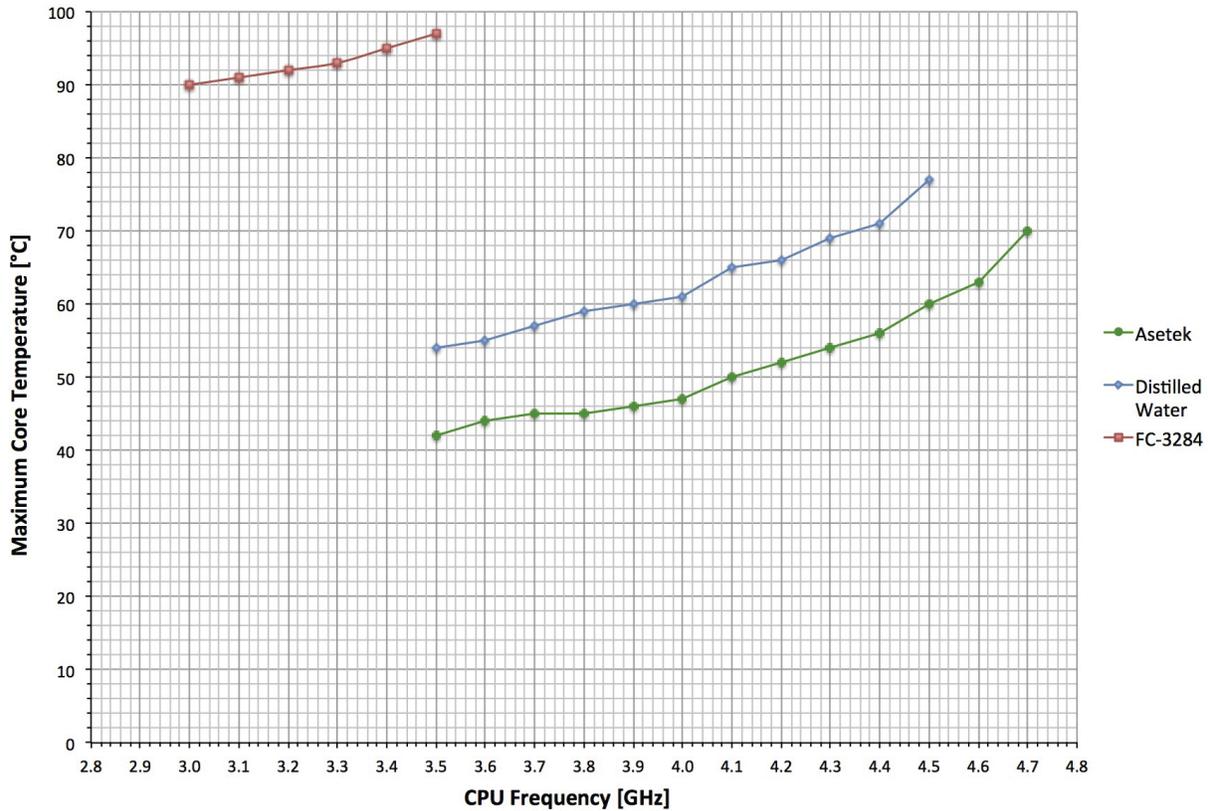


Figure 5–8: Maximum CPU core temperature comparison for the real CPU experiments using the Asetek, distilled water spray and FC-3284 spray cooling systems.

effectively a decreasing function of CPU frequency (i.e. cooling systems must maintain lower core temperatures at higher frequencies to pass the stress test).

The temperatures of the CPU package’s components are proportional to the electrical power input to the CPU, which is converted to heat and is transferred away by means of the cooling system. The power input to the CPU, as measured by sensors embedded in the server, as a function of CPU frequency for each of the tested cooling systems is plotted in figure 5–9. Similar to the core temperatures, the CPU power shows an exponentially increasing trend with increased frequency. This trend highlights one of the main difficulties in overclocking a processor, which is the exponentially increasing amount of energy that must be transferred away from the processor by an efficient cooling process. Furthermore, this indicates that a dramatic, rather than incremental, improvement in processor cooling is required to achieve substantially faster processing speeds.

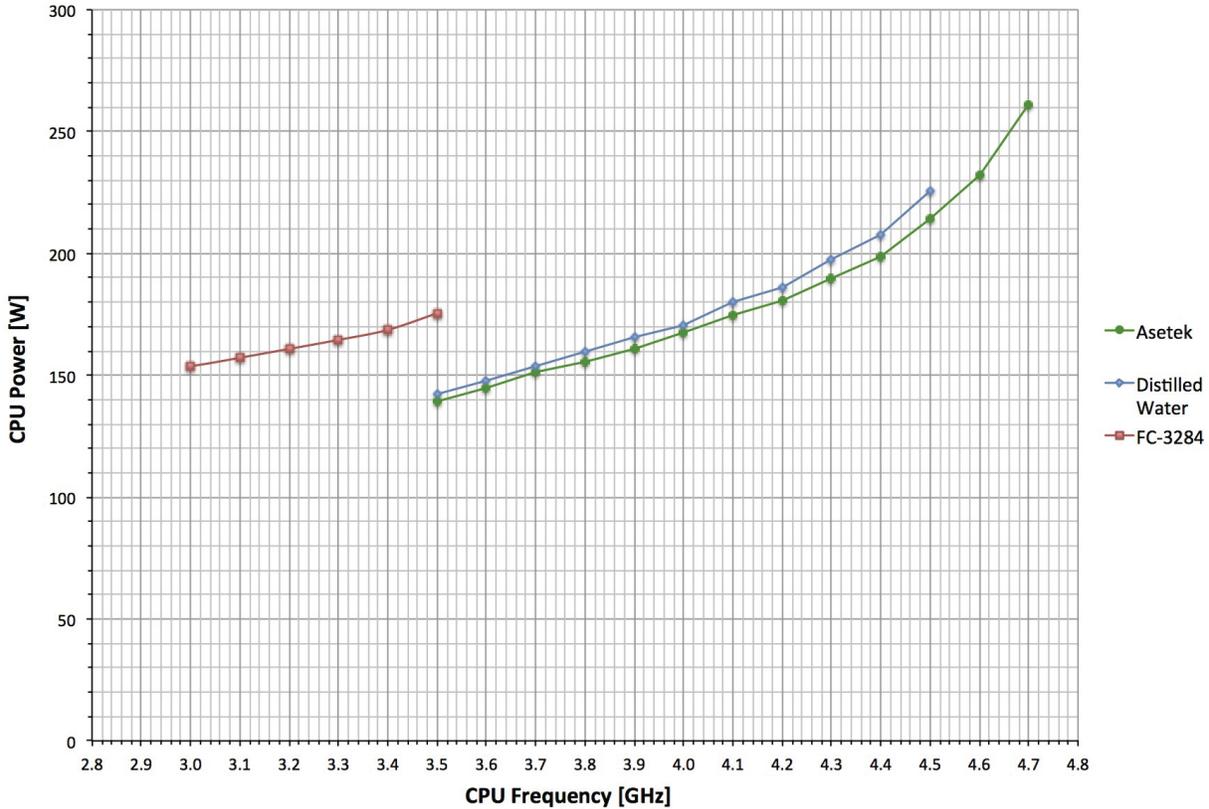


Figure 5–9: CPU power comparison for the real CPU experiments using the Asetek, distilled water spray, and FC-3284 spray cooling systems.

Comparing the three curves of figure 5–9, it is observed that the cooling systems induce different amounts of CPU power at the same frequency, despite all other settings remaining equal. The FC-3284 system induces significantly more CPU power at 3.5 GHz compared to the other systems, and the distilled water experiments show slightly higher power consumption compared to the Asetek experiments. This power difference is due to the differences in core temperatures. At elevated temperatures, such as in the case of the FC-3284 experiments, leakage power increases, which is due to leakage current between differently doped parts of the transistors (Kim *et al.* 2003). Therefore, improved processor cooling has the added benefit of reducing CPU power consumption.

Figure 5–10 plots the sensible heat transfer rate (calculated with equation 5.2) as a function of CPU frequency for the distilled water and FC-3284 experiments. There is no sensible heat transfer data for the Asetek system because its inlet and outlet temperatures of the

device could not be readily measured. One observes that the FC-3284 coolant temperature never reached its boiling point; therefore, no latent heat transfer occurred in those tests, and therefore both curves actually represent total heat transfer rate. These curves follow similar trends to the CPU power curve, which is expected because, ideally, most power consumed by a processor should be transferred away by the cooling system in the form of heat. However, a comparison of figures 5–9 and 5–10 reveals that there are large differences between the CPU power consumption and heat transfer rate in both cases. These differences are mainly due to heat losses through the CPU’s other thermal pathways, such as through its pins or its seat in the CPU socket. It appears that these losses were even greater in the FC-3284 case because the FC-3284 cooling system represented a pathway that was more thermally resistive relative to the distilled water system.

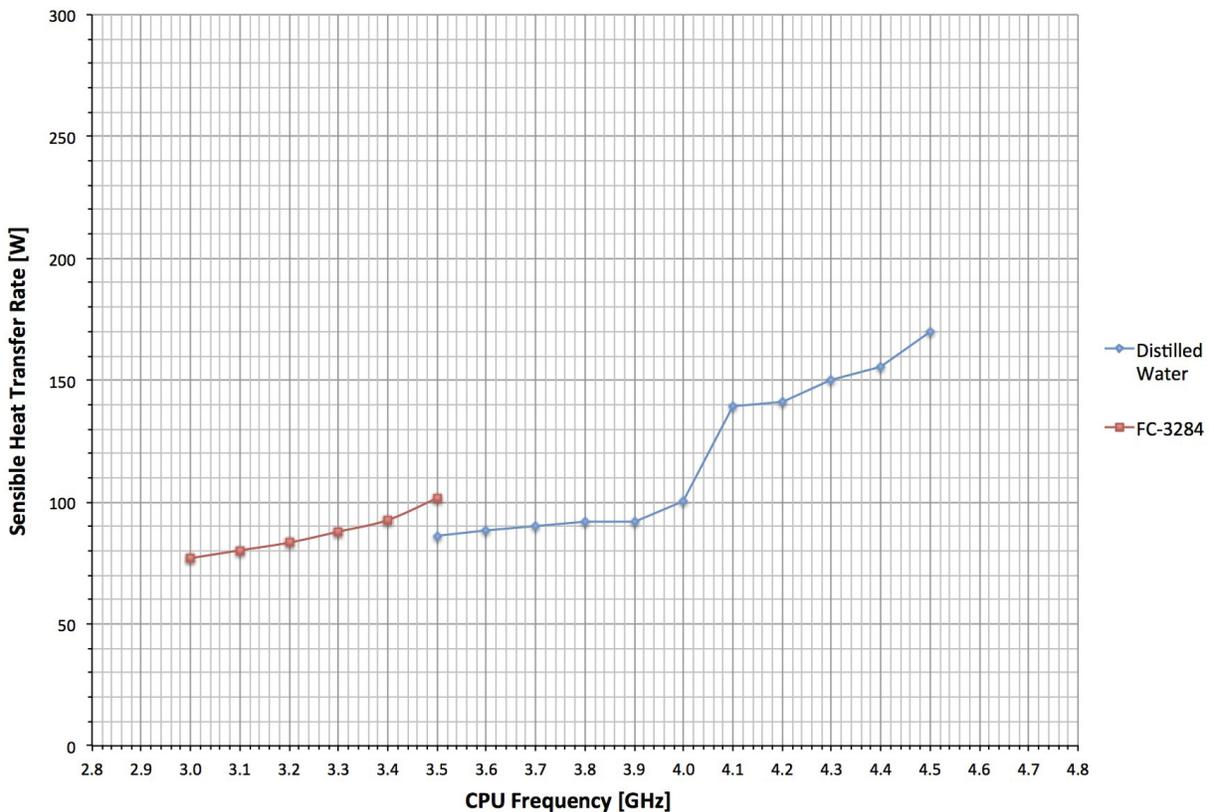


Figure 5–10: Sensible heat transfer rate comparison for the real CPU experiments using the distilled water, and FC-3284 spray cooling systems.

In figure 5–10, one observes an apparent discrepancy in the distilled water curve where there is an uncharacteristically large increase in heat transfer between frequencies of 4.0 GHz and 4.1 GHz. Moreover, there are noticeably larger than expected increases, at this point, in CPU power and core temperatures as well, although not nearly as pronounced. The fact that the increase in heat is not matched by an equal increase in CPU power indicates that, perhaps, the on-chip logic for CPU power distribution allocates a larger portion of the input power to the cores at frequencies above 4.0 GHz. This would move more power dissipation to the center of the CPU, where the resulting heat is more easily transferred to the coolant, and away from the outer components, such as the memory controller. However, this can only be hypothesized since the internal logic of the 5960X CPU is not available for review.

Further insights can be made by comparing the heat transfer results of the real CPU tests to the results of the simulated CPU tests. The highest sensible heat transfer rate, at 4.5 GHz, is 170 W for the real CPU experiments with distilled water. At a similar heat transfer rate of 172 W, during the simulated CPU experiments, the approximate surface temperature was 53.5 °C, which is well below Intel’s recommended maximum IHS surface temperature of 66.8 °C. When the CPU frequency was increased to 4.6 GHz and the distilled water cooling system failed, the heat transfer rate would likely not have increased past 200 W based on an extrapolation of the curve for distilled water. 200 W corresponds to an approximate surface temperature of 61 °C, which is still below 66.8 °C. Therefore, it is unlikely that the cooling system failed at 4.6 GHz due to an average surface temperature above the maximum allowable IHS surface temperature, so the failure is likely due to variations in temperature and heat transfer across the surface of the IHS. The majority of CPU heat generation occurs in the cores, which are located near the center of the die. The spray cooling system targets the surface fairly uniformly; whereas, the microchannel design of the Asetek system specifically targets the central area of the IHS, near the cores. Targeting the cores, or ‘hotspots,’ is an effective processor cooling strategy that is not employed in the current spray cooling design.

Future design iterations could benefit from either a spray nozzle array pattern or a specialty spray nozzle that specifically targets the processor cores.

## CHAPTER 6

### Conclusion

The concluding chapter of this thesis first reviews the work that was undertaken. Plans and recommendations for future work to extend the project are then presented.

#### 6.1 Review of the Thesis

The overall objective of the present work is to improve the performance of processors by advancing the field of processor cooling. The first sections of Chapter 1 provided background information regarding processor cooling that built motivation towards this goal. Although significant amounts of research have been conducted on high-heat-flux cooling methods, very few commercially-available processor cooling devices employ any type of advanced cooling technology. Therefore, it was determined that an experimental apparatus that is capable of testing novel processor cooling devices in a realistic testing environment would benefit the research and development of advanced processor cooling methods.

Subsequently, multiple cooling methods having the potential to improve processor cooling were identified in a review of high-heat-flux cooling methods. The majority of these methods, including microchannels, sprays and jet impingement, required the use of a coolant in a single-phase or two-phase cooling system. It was therefore determined that the experimental apparatus would require a fluid loop capable of handling both single-phase and two-phase fluid flow. Furthermore, this review resulted in the design of a spray cooling device as the first attempt at creating an improved processor cooling methodology.

The development of the experimental apparatus (covered in Chapter 2) began with the fluid loop, which went beyond a simple single-pathway closed loop, to include multiple redirection valves and a backup reservoir as additional safety mechanisms. Furthermore, five temperature sensors and three pressure sensors to monitor the physical properties of the

flow were incorporated in its design. Subsequently, a simulated CPU, consisting of a heated surface with multiple integrated sensors, was added to the apparatus. For the real CPU experiments, Intel's 5960X processor was selected, which was used with CIARA's HF320 server. The entire experimental apparatus was built on a wooden platform that also housed the control and data acquisition system. A Raspberry Pi computer was used for the master controller, and it ran a custom-built Python application that handled all of the control and data acquisition, while also providing a graphical user interface to facilitate experimentation.

To test the new experimental apparatus and validate its operation, a CPU cooling device was prototyped with the aim of improving processor cooling. The survey of cooling methods, in Chapter 1, led to a spray cooling approach, and a novel design for the spray nozzle was conceived in Chapter 3. In this design, the spray targeted the CPU's IHS directly, and the device was fashioned to be compatible with a standard CPU socket. Two coolants were selected to test in the new device: (i) distilled water as a single-phase coolant, and (ii) FC-3284 as a two-phase coolant.

Chapter 4 described the methods used to conduct the two different types of experiments. The simulated CPU experiments were performed first, and they involved testing both coolants using the spray cooling device over a range of flow and heat transfer rates. The real CPU experiments involved many of the same procedures as the simulated CPU experiments, but they additionally incorporated the computer system and its stress testing software. These tests were performed over a series of CPU frequencies rather than heat transfer rates, and only at the maximum flow rates of each coolant.

The simulated CPU experiments (described in Chapter 5) resulted in linear trends in the heat transfer rate as a function of surface temperature in the single-phase regime of both coolants. For surface temperatures above the boiling point of FC-3284, the tests proved that latent heat transfer considerably improves the heat transfer rate in the two-phase regime. However, the single-phase distilled water spray nevertheless managed to achieve a much higher heat transfer rate (214 W) at the main temperature of interest (66.8 °C) compared to

the two-phase FC-3284 spray (106 W). This result was partially due to the higher specific heat capacity of water and partially due to the under utilization of the two-phase spray, which were highlighted by the vapour quality and spray efficiency plots. Because the spray only significantly vaporized at surface temperatures above 78 °C, the advantages of two-phase cooling were not being fully utilized in the required temperature range. Overall, the simulated CPU experiments indicated that the distilled water spray was the more promising candidate for processor cooling.

Before beginning the real CPU experiments with the spray cooling device, a performance benchmark was determined using the Asetek cooling system that is currently implemented in the HF320 server. The final pass/fail results of the tests confirmed that FC-3284 was not a suitable coolant for this application, and revealed that the distilled water spray, in its current configuration, was not capable of reaching the same performance levels as the Asetek system (maximum CPU frequency of 4.5 GHz versus 4.7 GHz). A more in-depth analysis of the data underscored the current limitations of the distilled water spray cooling device. The cores were 10–15 °C hotter on average relative to the Asetek results. In addition to causing operational error in the CPU, the elevated temperatures actually increased the amount of power required to run the CPU, thus compounding the problem. Therefore, future design iterations of the spray cooling device and new devices altogether must maintain the CPU temperatures much lower and, perhaps, target the cores more directly to achieve an improvement in processor cooling.

## **6.2 Extensions of the Present Work**

There are innumerable ways to extend the present work because the experimental apparatus provides a platform for quickly testing any new processor cooling technologies. Eventually, other cooling methods, such as microchannels and jet impingement, will be incorporated into new processor cooling device designs and tested using the apparatus. However, the testing of new configurations with the existing spray cooling device is the primary objective of

future work. Secondary extension objectives include improving the experimental apparatus, giving it additional capabilities.

There remain numerous ways to configure the current spray cooling device design to result in improved processor cooling. The first improvement would be to permit higher flow rates by using a spray nozzle with a slightly larger orifice diameter. Higher flow rates could also be achieved with an array of spray nozzles. Furthermore, an array of nozzles could be arranged to target the cores of the processor. As mentioned in section 5.2.2, new coolants with lower boiling points, such as Novec 7000, could also make two-phase processor spray cooling at atmospheric pressure feasible and succeed where FC-3284 could not.

The first improvement to be made to the experimental apparatus is to add additional temperature sensors to the copper block heated surface of the simulated CPU. Currently, the copper block contains one sensor to approximate the surface temperature. If two planes of five-sensor arrays at different depths were incorporated into the block, not only would a surface temperature profile be observable, but an enhanced estimation of heat transfer rates by conduction to the surface could also be made. This would eliminate the need to predict heat transfer rates in two-phase cooling. Subsequently, adding a device to the fluid loop to subcool the coolant prior to it entering the cooling device would add further control to the experiments. Since vapour-compression refrigeration systems are not feasible in the small volume of a computer case, a more innovative method, such as thermoelectric cooling, could be tested to perform the subcooling instead.

## REFERENCES

- Betz, A. R. *et al.* (2011). “Significant boiling enhancement with surfaces combining superhydrophilic and superhydrophobic patterns”. *Micro Electro Mechanical Systems (MEMS), 2011 IEEE 24th International Conference on*. IEEE, pp. 1193–1196.
- Bogojevic, D. (2011). “Experimental investigation of non-uniform heating effect on flow boiling instabilities in a microchannel-based heat sink”. *International Journal of Thermal Sciences* 50.3, pp. 309–324.
- Bostanci, H. *et al.* (2009). “Spray cooling of power electronics using high temperature coolant and enhanced surface”. *2009 IEEE Vehicle Power and Propulsion Conference*, pp. 609–613.
- Brunschwiler, T. *et al.* (2008). “Forced convective interlayer cooling in vertically integrated packages”. *Thermal and Thermomechanical Phenomena in Electronic Systems, 2008. ITherm 2008. 11th Intersociety Conference on*. IEEE, pp. 1114–1125.
- Chowdhury, I. (2009). “On-chip cooling by superlattice-based thin-film thermoelectrics”. *Nature Nanotechnology* 4.4, p. 235.
- Ebadian, M. A. and Lin, C. X. (2011). “A Review of High-Heat-Flux Heat Removal Technologies”. *Journal of Heat Transfer-Transactions of the Asme* 133.11.
- Fabbri, M., Jiang, S., and Dhir, V. K. (2003). “Comparative Study of Spray and Multiple Micro Jets Cooling for High Power Density Electronic Applications”. *ASME 2003 International Mechanical Engineering Congress and Exposition*. American Society of Mechanical Engineers, pp. 657–665.
- Garimella, S. V. (2008). “Thermal challenges in next-generation electronic systems”. *IEEE Transactions on Components and Packaging Technologies* 31.4, pp. 801–815.

- Gong, L. (2011). “Parametric numerical study of flow and heat transfer in microchannels with wavy walls”. *Journal of Heat Transfer* 133.5.
- Hassan, M. I. (2015). “An experimental study of heat pipe performance using nanofluids”. *International Journal of Green Energy* 12.3, pp. 225–229.
- Haywood, A. M. (2015). “The relationship among CPU utilization, temperature, and thermal power for waste heat utilization”. *Energy Conversion and Management* 95, pp. 297–303.
- Hinton, M. J. and Mydlarski, L. B. (2016). “CPU Cooling System with Direct Spray Cooling”. U.S. Non-Provisional Patent, Pending.
- Hirshfeld, H. (2006). “High heat flux cooling of accelerator targets with micro-channels”. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 562.2, pp. 903–905.
- Hu, X. *et al.* (2005). “Thermal characterization of vertically-oriented carbon nanotubes on silicon”. *Semiconductor Thermal Measurement and Management Symposium, 2005 IEEE Twenty First Annual IEEE*. IEEE, pp. 292–297.
- Kim, N. S. *et al.* (2003). “Leakage current: Moore’s law meets static power”. *Computer* 36.12, pp. 68–75.
- Koncar, B. (2010). “Effect of nozzle sizes on jet impingement heat transfer in He-cooled divertor”. *Applied Thermal Engineering* 30.6-7, pp. 697–705.
- Koo, J.-M. *et al.* (2002). “Convective boiling in microchannel heat sinks with spatially-varying heat generation”. *Thermal and Thermomechanical Phenomena in Electronic Systems, 2002. IThERM 2002. The Eighth Intersociety Conference on*. IEEE, pp. 341–346.
- Koo, J.-M. *et al.* (2005). “Integrated microchannel cooling for three-dimensional electronic circuit architectures”. *Journal of heat transfer* 127.1, pp. 49–58.
- Lee, J. (2009). “Critical heat flux for subcooled flow boiling in micro-channel heat sinks”. *International Journal of Heat and Mass Transfer* 52.13-14, pp. 3341–3352.

- Lee, J. H. (2008). “Effective viscosities and thermal conductivities of aqueous nanofluids containing low volume concentrations of Al<sub>2</sub>O<sub>3</sub> nanoparticles”. *International Journal of Heat and Mass Transfer* 51.11-12, pp. 2651–2656.
- Lienhard, J. H. (1973). “On the commonality of equations for natural convection from immersed bodies”. *International Journal of Heat and Mass Transfer* 16.11, pp. 2121–2123.
- Liu, X. and Lienhard, J. H. (1992). “Extremely high heat flux removal by subcooled liquid jet impingement”. *ASME HEAT TRANSFER DIV PUBL HTD., ASME, NEW YORK, NY(USA), 1992* 217, pp. 11–20.
- Liu, Z. (2006). “Critical heat flux of steady boiling for water jet impingement in flat stagnation zone on superhydrophilic surface”. *Journal of Heat Transfer* 128.7, pp. 726–729.
- Mahajan, R., Chiu, C. P., and Chrysler, G. (2006). “Cooling a microprocessor chip”. *Proceedings of the IEEE* 94.8, pp. 1476–1485.
- Marcinichen, J. B. (2010). “Cooling of microprocessors with micro-evaporation: A novel two-phase cooling cycle”. *International Journal of Refrigeration* 33.7, pp. 1264–1276.
- Mochizuki, M. *et al.* (2009). “Thermal management in high performance computers by use of heat Pipes and vapor chambers, and the challenges of global warming and environment”. *Microsystems, Packaging, Assembly and Circuits Technology Conference, 2009. IMPACT 2009. 4th International*. IEEE, pp. 191–194.
- Morton, J. J. *et al.* (2011). “Embracing the quantum limit in silicon computing”. *Nature* 479.7373, pp. 345–353.
- Mudawar, I. (1999). “Ultra-high critical heat flux (CHF) for subcooled water flow boiling-I: CHF data and parametric effects for small diameter tubes”. *International Journal of Heat and Mass Transfer* 42.8, pp. 1405–1428.
- Murshed, S. M. S. (2009). “A combined model for the effective thermal conductivity of nanofluids”. *Applied Thermal Engineering* 29.11-12, pp. 2477–2483.
- Nguyen, C. T. (2007). “Heat transfer enhancement using Al<sub>2</sub>O<sub>3</sub>-water nanofluid for an electronic liquid cooling system”. *Applied Thermal Engineering* 27.8-9, pp. 1501–1506.

- Nguyen, T. *et al.* (2000). “Use of heat pipe/heat sink for thermal management of high performance CPUs”. *Sixteenth Annual IEEE Semiconductor Thermal Measurement and Management Symposium (Cat. No.00CH37068)*, pp. 76–79.
- Pais, M. R. (1992). “Surface roughness and its effects on the heat transfer mechanism in spray cooling”. *Journal of Heat Transfer* 114.1, pp. 211–219.
- Peterson, G. P. (1990). “Thermal Control of Electronic Equipment and Devices”. *Advances in Heat Transfer* 20.C, pp. 181–314.
- Prasher, R. (2006). “Thermal interface materials: Historical perspective, status, and future directions”. *Proceedings of the IEEE* 94.8, pp. 1571–1586.
- Saidur, R. (2011). “A review on applications and challenges of nanofluids”. *Renewable and Sustainable Energy Reviews* 15.3, pp. 1646–1668.
- Sharma, A. *et al.* (2009). “Review on thermal energy storage with phase change materials and applications”. *Renewable and Sustainable Energy Reviews* 13.2, pp. 318–345.
- Sharma, C. S. (2015). “A novel method of energy efficient hotspot-targeted embedded liquid cooling for electronics: An experimental study”. *International Journal of Heat and Mass Transfer* 88, pp. 684–694.
- Shehabi, A *et al.* (2016). “United States data center energy usage report”. *Lawrence Berkeley National Laboratory, Berkeley, California. LBNL-1005775 Page 4*.
- Silk, E. A. (2006). “Spray cooling of enhanced surfaces: Impact of structured surface geometry and spray axis inclination”. *International Journal of Heat and Mass Transfer* 49.25-26, pp. 4910–4920.
- Silverman, I. (2006). “High heat-flux accelerator targets: Cooling with liquid metal jet impingement”. *International Journal of Heat and Mass Transfer* 49.17-18, pp. 2782–2792.
- Srikar, R. *et al.* (2009). “Nanofiber coating of surfaces for intensification of drop or spray impact cooling”. *International Journal of Heat and Mass Transfer* 52.2526, pp. 5814–5826.

- Steinhart, J. S. and Hart, S. R. (1968). “Calibration curves for thermistors”. *Deep Sea Research and Oceanographic Abstracts* 15.4, pp. 497–503.
- Sung, M. K. and Mudawar, I. (2009). “Single-Phase and Two-Phase Hybrid Cooling Schemes for High-Heat-Flux Thermal Management of Defense Electronics”. *Journal of Electronic Packaging* 131.2, pp. 021013–021013.
- Takeuchi, I. and Sandeman, K. (2015). “Solid-state cooling with caloric materials”. *Physics Today* 68.12, pp. 48–54.
- Tavoularis, S. (2005). *Measurement in fluid mechanics*. Cambridge University Press.
- Tschudi, W. *et al.* (2003). “Data Centers and Energy Use Lets Look at the Data”. *Proceedings of the 2003 ACEEE Summer Study on Energy Efficiency in Industry*. Rye Brook, NY: American Council for an Energy Efficient Economy, in Washington, DC. July.
- Tuckerman, D. B. and Pease, R. F. W. (1981). “High-Performance Heat Sinking for VLSI”. *IEEE Electron Device Letters* EDL-2.5.
- Verma, K *et al.* (1998). “Real-time warpage measurement of electronic components with variable sensitivity”. *Electronic Components & Technology Conference, 1998. 48th IEEE*. IEEE, pp. 975–980.

## APPENDIX A: Heat Transfer Rate Estimation

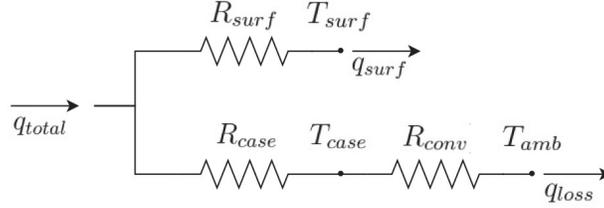


Figure A–1: Thermal resistance diagram for heat transfer model.

Using the principles of heat transfer and the notion of thermal resistance, a simple model was developed to estimate the heat transfer rate across the surface of the simulated CPU. The following two equations, which are also represented graphically in figure A–1, were used to govern the model:

$$q_{total} = q_{surf} + q_{loss}, \text{ and} \quad (\text{A.1})$$

$$q_{surf}R_{surf} + T_{surf} = q_{loss}(R_{case} + R_{conv}) + T_{amb}, \quad (\text{A.2})$$

where:

- $q_{surf}$  is the heat transfer rate across the surface (denoted simply as  $q$  in Chapter 5), which is the parameter of interest;
- $q_{total}$  is the total heat generated by the simulated CPU, which is equal to the simulated CPU power;
- $q_{loss}$  is the heat loss;
- $T_{surf}$  is the surface temperature;
- $T_{case}$  is the case temperature;
- $T_{amb}$  is the ambient temperature;
- $R_{surf}$  is the thermal resistance between the heat source and the surface;
- $R_{case}$  is the thermal resistance between the heat source and the case; and

- $R_{conv}$  is the thermal resistance to natural convection from the case to the environment.

The second equation assumes that radiation is negligible, and heat transfer from the case to the environment is solely due to natural convection. A few more assumptions need to be made to use the above equations in the model. Steady state is assumed, which is a reasonable assumption because the measurements are only made after steady state is achieved. Next, the thermal resistances  $R_{surf}$  and  $R_{case}$  are assumed to be constant. Subsequently, the thermal resistance to natural convection is assumed to be of form:  $R_{conv} = C_{conv}(T_{case} - T_{amb})^{-\frac{1}{4}}$ , where  $C_{conv}$  is an experimentally determined constant. This model form is derived assuming natural convection of an irregular solid (Lienhard 1973). Finally,  $T_{surf}$  and  $T_{case}$  are assumed to be uniform across their respective surfaces.

The governing equations were then combined into the following representation:

$$q_{surf} = \frac{(q_{total} - q_{surf})(R_{case} + C_{conv}(T_{case} - T_{amb})^{-\frac{1}{4}}) - (T_{surf} - T_{amb})}{R_{surf}}. \quad (\text{A.3})$$

Using the data from the distilled water experiments, a nonlinear model fit in Mathematica was used to find the unknowns and yielded:  $R_{surf} = 1.373 \text{ K/W}$ ,  $R_{case} = 1.995 \text{ K/W}$ , and  $C_{conv} = 4.112 \text{ K}^{\frac{5}{4}}/\text{W}$ .

The final model equation for the heat transfer rate across the surface is:

$$q_{surf} = \frac{q_{total}(R_{case} + C_{conv}(T_{case} - T_{amb})^{-\frac{1}{4}}) - (T_{surf} - T_{amb})}{R_{surf} + R_{case} + C_{conv}(T_{case} - T_{amb})^{-\frac{1}{4}}}. \quad (\text{A.4})$$

The distilled water data was then used to determine how well the model fit, and an average percent difference of 2.23% was found, and the standard deviation was 2.94 W. This is precise enough for the heat transfer estimation of the two-phase cooling experiments.

## APPENDIX B: Uncertainty Analysis

The present appendix details the uncertainty analysis performed for the reported measurements. The analysis follows the uncertainty theory outlined by Tavoularis (2005), which employs a 95% confidence interval. Generally, for each property of interest, there is an analytical expression  $y = y(x_1, x_2, \dots, x_M)$  that is used to evaluate the property based on  $M$  subsidiary properties,  $x_m$ . The property can be any measured property, such as temperature or pressure, or a derived property, such as heat transfer rate or spray efficiency. The subsidiary properties are the properties, such as flow rate, that are required to evaluate the desired property. The uncertainty  $u_y$  of property  $y$  is calculated from the bias limit  $b_y$  and the precision limit  $p_y$  as:

$$u_y = \sqrt{b_y^2 + p_y^2}. \quad (\text{B.1})$$

Bias limits indicate that the experimenter is 95% confident that the true bias error is less than the reported bias limit. The bias limit for measured properties that have  $K$  different components or steps of the measuring process is determined as follows:

$$b_y = \sqrt{\sum_k^K b_k^2}, \quad (\text{B.2})$$

where  $b_k$  are the bias limits for each part of the measuring process. For derived properties, the bias limit, when all subsidiary properties are measured independently, is computed using:

$$b_y = \sqrt{\sum_m^M \left( \frac{\partial y}{\partial x_m} b_{x_m} \right)^2}, \quad (\text{B.3})$$

where  $b_{x_m}$  is the bias of subsidiary property  $x_m$ . One bias that is common to all measured properties is the quantization uncertainty, which Tavoularis approximates as one half of the resolution of the analogue-to-digital converter. Analogue-to-digital conversion is handled by the Arduinos, which perform 10-bit discretization over a voltage range of 5 V; therefore, the

bias limit for quantization uncertainty is 2.4 mV.

The precision limit is the uncertainty due to random error. It can be estimated based on statistical properties of the measurement. All reported values are measurement averages over a series of  $N$  individual measurements  $y_n$ . (Herein,  $N$  is slightly different for each measurement, but it is typically well over 100 samples.) The statistical properties needed to calculate the precision limit are the mean  $\mu_{yN}$ , which is also the reported value, and the variance  $\sigma_{yN}^2$ :

$$\mu_{yN} = \frac{1}{N} \sum_n^N y_n, \text{ and} \quad (\text{B.4})$$

$$\sigma_{yN}^2 = \frac{1}{N-1} \sum_n^N (y_n - \mu_{yN})^2. \quad (\text{B.5})$$

Then, the precision can be computed with:

$$p_y = \frac{2\sigma_{yN}}{\sqrt{N}}. \quad (\text{B.6})$$

The remainder of this appendix focuses on each reported measurement type, and determines their overall uncertainty. First, the measured temperature and pressure are analyzed. Then, derived properties analyses, such as heat transfer rate, vapour quality and spray efficiency, follow.

### B.1 Simulated CPU Surface Temperature

The surface temperature of the simulated CPU was measured using an RTD sensor that has a measurement range of -4.0–350.0 °F (-20.0–176.7.0 °C) and an accuracy of  $\pm 0.12\%$  of full scale. The RTD probe is connected to a 4–20 mA transmitter with an output range of 0.0–500.0 °F (-17.8–260.0 °C) and accuracy of 0.10%. The transmitter current is converted to a voltage across a  $250.0 \pm 2.5 \Omega$  load resistor. The maximum surface temperature was 92.5 °C, so the maximum bias limit due to inaccuracy in the load resistance value is 1.8 °C. The Arduino Mega discretizes the analogue voltage with a resolution of 4.8 mV or 0.34 °C.

Table B–1: Summary of bias errors for the simulated CPU surface temperature measurements.

Bias Error	$b_k$
RTD probe	0.24 °C
RTD transmitter	0.28 °C
Load resistor	1.8 °C
Quantization	0.17 °C

These bias errors are summarized in table B–1 along with their corresponding bias limits. The resulting bias limit is  $b_{T_{surf}} = 1.8$  °C, which is clearly dominated by the load resistor bias.

The precision limit was calculated for each series of surface temperature measurements, and the values range from 0.011 to 0.15 °C. As a conservative measure, the precision limit was set to 0.15 °C; however, this largest error still does not affect the overall uncertainty relative to the much larger bias limit. The final uncertainty for surface temperature measurements is therefore:

$$u_{T_{surf}} = 1.8 \text{ } ^\circ\text{C}.$$

## B.2 Simulated CPU Case Temperature

The case temperature of the simulated CPU was measured using a type J thermocouple sensor that has a measurement range of 32–1380 °F (0.0–748.9 °C) and an accuracy of  $\pm 0.75\%$  of full scale. Its 4–20 mA transmitter has an output range of 32.0–1292.0 °F (0.0–700.0 °C) and accuracy of  $\pm 0.10\%$ . Like the RTD circuit, the transmitter current is read by the Arduino as a voltage with the use of a  $250.0 \pm 2.5$   $\Omega$  load resistor. The maximum bias limit due to inaccuracy in the load resistance value, with a maximum case temperature of 249.7 °C, is 4.2 °C. The analogue voltage is discretized with a resolution of 4.8 mV or 0.86 °C. Table B–2 provides the corresponding bias limits to these bias errors. The overall bias limit is  $b_{T_{case}} = 7.1$  °C. The precision limits ranged from 0.07 to 0.22 °C, and the high end of the range was selected again as the representative precision limit. The final uncertainty for case

Table B-2: Summary of bias errors for the simulated CPU case temperature measurements.

Bias Error	$b_k$
Thermocouple probe	6 °C
Thermocouple transmitter	0.7 °C
Load resistor	4 °C
Quantization	0.4 °C

temperature measurements is:

$$u_{T_{case}} = 7^{\circ}C.$$

### B.3 CPU Cooling Device Inlet and Outlet Temperatures

The temperatures of the flow entering and exiting the CPU cooling device were sensed by XSPC inline temperature sensors, which are 10 k $\Omega$  thermistors. As mentioned in section 2.1.6, the sensors had not been pre-calibrated, so a calibration was performed to determine the coefficients of the Steinhart-Hart model equation (Steinhart and Hart 1968) that maps the temperature to the resistance output. The Steinhart-Hart equation is:

$$T = [A + B \ln(R) + C[\ln(R)]^3]^{-1}, \quad (\text{B.7})$$

where  $T$  is the temperature in K,  $R$  is the resistance in  $\Omega$  across the thermistor, and the coefficients were determined to be  $A = 1.21 \times 10^{-3}$ ,  $B = 1.89 \times 10^{-4}$  and  $C = 5.22 \times 10^{-7}$ . This calibration process used a digital multimeter (MS8268) to measure resistance ( $\pm 1.2\%$  of reading + 2 digits), and a NIST-calibrated digital thermometer with an accuracy of  $\pm 0.5^{\circ}C$ . The recorded measurements from these digital devices were averaged over a sufficient number of measurements such that the precision limit was assumed to be negligible. Without a temperature test chamber available, the sensors could only be calibrated with the three points necessary to solve for the Steinhart-Hart coefficients. Therefore, there were not enough additional data points to determine the exact uncertainty  $u_{xspc}$  of the model fit equation. Instead, the uncertainty was approximated by:

$$u_{xspc} = \sqrt{b_{therm}^2 + \left(\frac{\partial T}{\partial R} b_{mult}\right)^2}, \quad (\text{B.8})$$

where  $b_{therm} = 0.5 \text{ }^\circ\text{C}$  is the bias limit of the thermometer, and  $b_{mult} = 0.012R$  is the bias limit of the multimeter. The uncertainty is a function of the resistance, and it can be approximated by a decreasing linear function over the resistance range of interest (2,530–10,500  $\Omega$ , which is based on the full temperature range of measurements). The uncertainty of the XSPC sensor is thus approximated as:

$$u_{xspc} = 0.642 \text{ }^\circ\text{C} - \left( 3.73 \times 10^{-6} \frac{\text{ }^\circ\text{C}}{\Omega} \right) R. \quad (\text{B.9})$$

The thermistor was then connected in series with a  $10,000 \pm 100 \text{ } \Omega$  resistor to create a voltage divider circuit, and the voltage was read by the 10-bit analogue-to-digital converter of the Arduino. The final bias limit was computed as:

$$b_T = \sqrt{u_{xspc}^2 + \left( \frac{\partial T}{\partial R} \frac{\partial R}{\partial R_{10k}} b_{R_{10k}} \right)^2 + \left( \frac{\partial T}{\partial R} \frac{\partial R}{\partial div} b_{div} \right)^2}, \quad (\text{B.10})$$

where  $\frac{\partial T}{\partial R} \frac{\partial R}{\partial R_{10k}} b_{R_{10k}}$  is the bias due to the 10 k $\Omega$  resistor, and  $\frac{\partial T}{\partial R} \frac{\partial R}{\partial div} b_{div}$  is the bias due to quantization error. These two biases can be rewritten as:

$$\frac{\partial T}{\partial R} \frac{\partial R}{\partial R_{10k}} b_{R_{10k}} = 0.01 \frac{\partial T}{\partial R}, \text{ and} \quad (\text{B.11})$$

$$\frac{\partial T}{\partial R} \frac{\partial R}{\partial div} b_{div} = \frac{1}{2,046} R \left( \frac{R}{10,000 \text{ } \Omega} + 1 \right) \frac{\partial T}{\partial R}. \quad (\text{B.12})$$

A plot of  $b_T$  reveals that the bias limit is linear and slightly decreasing with  $R$  on the resistance range 2,530–10,500  $\Omega$ . The maximum bias is  $0.71 \text{ }^\circ\text{C}$ , and the minimum is  $0.66 \text{ }^\circ\text{C}$ ; therefore, the bias limit should be simplified to a constant value of  $0.7 \text{ }^\circ\text{C}$ . Furthermore, the precision limits of both the inlet and outlet temperatures are negligible relative to the bias. Thus, the overall uncertainty of these inline flow temperature measurements is:

$$u_{T_{in}} = u_{T_{out}} = 0.7 \text{ }^\circ\text{C}.$$

Table B-3: Summary of bias errors for the ambient temperature measurements.

Bias Error	$b_k$
DHT temperature sensor	0.5 °C
Quantization	0.0009 °C

#### B.4 Ambient Temperature

The ambient temperature sensor (DHT22/AM2302) has a temperature range of -40.0–80.0 °C and an accuracy of  $\pm 0.5$  °C. Furthermore, it transmits the temperature measurement by way of a 16-bit digital signal. The resulting biases are summarized in table B-3, and the maximum precision limit is 0.005 °C. The DHT accuracy is the only significant factor in the uncertainty; thus, the uncertainty of the ambient temperature measurements is:

$$u_{T_{amb}} = 0.5 \text{ } ^\circ\text{C}.$$

#### B.5 Pressure

The pressure transmitters connected to the fluid loop have a range of 0–10 bar. Their accuracy is  $\pm 0.5\%$  of full scale (0.05 bar). They transmit a 4–20 mA signal that is converted to a voltage across a  $250 \pm 2.5$   $\Omega$  resistor, and that voltage is then read by the Arduino. All biases related to the pressure measurements are listed in table B-4. The largest precision limit was 0.009 bar, which brings the final pressure uncertainty to:

$$u_p = 0.07 \text{ bar}.$$

Table B-4: Summary of bias errors for the pressure transmitter measurements.

Bias Error	$b_k$
Pressure transmitter	0.05 bar
Load resistor	0.05 bar
Quantization	0.006 bar

## B.6 Sensible Heat Transfer Rate

The sensible heat transfer rate is a derived measurement that is computed as follows:

$$q_{sens} = \rho_l c_{p,l} Q (T_{out} - T_{in}). \quad (\text{B.13})$$

Assuming that the properties  $\rho_l$  and  $c_{p,l}$  are relatively certain and constant, the bias limit for sensible heat transfer rate can be calculated as:

$$b_{q_{sens}} = \rho_l c_{p,l} \sqrt{(T_{out} - T_{in})^2 b_Q^2 + Q^2 b_{T_{in}}^2 + Q^2 b_{T_{out}}^2}, \quad (\text{B.14})$$

where  $b_Q = 0.01Q$  is the flow rate bias limit. The bias limit is thus directly proportional to the flow rate. The average bias is  $b_{q_{sens}} = 10 \text{ W}$ . The precision limit is orders of magnitude below this, so it is not considered in the total uncertainty, which can be calculated using:

$$u_{q_{sens}} = \rho_l c_{p,l} Q \sqrt{0.0001(T_{out} - T_{in})^2 + 0.98^\circ C^2}. \quad (\text{B.15})$$

## B.7 Simulated CPU Power

The uncertainty in the simulated CPU power measurements (denoted with  $q_{total}$  rather than  $P_{ave}$  for consistency with appendices) was determined by analyzing the measurement error over a range of known power values. The measured power was compared to measurements with a KAW, which were made by monitoring its display over the measurement period. Ten measurements were made over a range from 48 W to 266 W. An analysis of the errors with respect to the KAW measurements revealed that the errors increased with higher powers. Consequently, the bias was determined to be  $0.006q_{total}$  or  $\pm 0.6\%$  of reading. The bias of the KAW readings is  $\pm 0.5\%$  of reading; thus, the bias limit is computed to be  $0.008q_{total}$  using equation B.2. Finally, the precision limits were determined to be negligible relative to the bias limits. Therefore, the uncertainty of the simulated CPU power readings is:

$$u_{q_{total}} = 0.008q_{total}. \quad (\text{B.16})$$

## B.8 Estimated Heat Transfer Rate

The estimated heat transfer rate for the two-phase tests, which is calculated using equation A.4 for  $q_{surf}$  from Appendix A, involves uncertainty from multiple sources. The uncertainty is calculated as follows:

$$u_{q_{surf}} = \sqrt{\sum_i^I \left( \frac{\partial q_{surf}}{\partial x_i} u_{x_i} \right)^2 + u_{fit}^2 + u_{q_{sens}}^2}, \quad (\text{B.17})$$

where  $u_{fit}$  is the uncertainty due to the errors in fit of the heat transfer rate estimates with the sensible heat transfer rate measurements,  $u_{q_{sens}}$  is the uncertainty in the sensible heat transfer rate measurements, and  $\sqrt{\sum_i^I \left( \frac{\partial q_{surf}}{\partial x_i} u_{x_i} \right)^2}$  is the uncertainty due to the parameters of the estimation equation, which expands to:

$$\begin{aligned} \sum_i^I \left( \frac{\partial q_{surf}}{\partial x_i} u_{x_i} \right)^2 &= \left( \frac{\partial q_{surf}}{\partial q_{total}} u_{q_{total}} \right)^2 + \left( \frac{\partial q_{surf}}{\partial T_{case}} u_{T_{case}} \right)^2 \\ &\quad + \left( \frac{\partial q_{surf}}{\partial T_{surf}} u_{T_{surf}} \right)^2 + \left( \frac{\partial q_{surf}}{\partial T_{amb}} u_{T_{amb}} \right)^2. \end{aligned} \quad (\text{B.18})$$

The uncertainties, represented by  $\sqrt{\sum_i^I \left( \frac{\partial q_{surf}}{\partial x_i} u_{x_i} \right)^2}$ , for each heat transfer estimate measurement were evaluated, and it was determined that they were negligible relative to the remaining uncertainty. The uncertainty,  $u_{fit}$ , was approximated by multiplying the standard deviation of the errors in fit (determined in Appendix A to be 2.94 W) by two. Finally, the uncertainty,  $u_{q_{sens}}$ , was approximated to 10 W by taking the average of sensible heat transfer rate measurement uncertainties, which were calculated using equation B.15. The total uncertainty of measurements using the estimated heat transfer rate equation is approximately:

$$u_{q_{surf}} = 12 \text{ W}.$$

## B.9 Vapour Quality

The vapour quality of the coolant was approximated using equation 5.4. The uncertainties of these derived measurements are calculated using:

$$u_\chi = \sqrt{\left(\frac{\partial\chi}{\partial q_{surf}}u_{q_{surf}}\right)^2 + \left(\frac{\partial\chi}{\partial q_{sens}}u_{q_{sens}}\right)^2 + \left(\frac{\partial\chi}{\partial Q}u_Q\right)^2}, \quad (\text{B.19})$$

which simplifies to:

$$u_\chi = \frac{1}{\rho_l \lambda Q} \sqrt{(12 W)^2 + (\rho_l c_{p,l} Q)^2 (0.0001 (T_{out} - T_{in})^2 + 0.98^\circ C^2) + 0.0001 (q_{surf} - q_{sens})^2}. \quad (\text{B.20})$$

An analysis of the resulting experimental uncertainties using this equation revealed that the uncertainty can reliably be rewritten as a simple function of the flow rate (for the FC-3284 coolant) as follows:

$$u_\chi = \frac{7 \times 10^{-8} \frac{m^3}{s}}{Q}. \quad (\text{B.21})$$

## B.10 Spray Efficiency

The uncertainty of a spray efficiency (equation 5.5) derived measurement is computed using:

$$u_\eta = \sqrt{\left(\frac{\partial\eta}{\partial q_{surf}}u_{q_{surf}}\right)^2 + \left(\frac{\partial\eta}{\partial Q}u_Q\right)^2 + \left(\frac{\partial\eta}{\partial T_{in}}u_{T_{in}}\right)^2}, \quad (\text{B.22})$$

which simplifies to:

$$u_\eta = \frac{1}{\rho_l Q (c_{p,l} (T_{sat} - T_{in}) + \lambda)} \sqrt{(12 W)^2 + (0.01 q_{surf})^2 + \left(\frac{(0.7 K) c_{p,l} q_{surf}}{c_{p,l} (T_{sat} - T_{in}) + \lambda}\right)^2}. \quad (\text{B.23})$$

Similar to the uncertainly analysis for vapour quality, an analysis of the uncertainties revealed that the spray efficiency uncertainty for the FC-3284 experiments can be further simplified

to a function solely dependent on flow rate:

$$u_\eta = \frac{5 \times 10^{-8} \frac{m^3}{s}}{Q}. \quad (\text{B.24})$$