Summer 2024

Machine Learning-Based Methodology for Multi-Objective and Multi-Design Variable Optimization of Finned Heat Sinks and Evaluation of Electrochemical Additive Manufactured Heat Sink Designs for Single-Phase Immersion Cooling

Joseph F. Herring II
University of Texas at Arlington

Machine Learning-Based Methodology for Multi-Objective and Multi-Design Variable Optimization of Finned Heat Sinks and Evaluation of Electrochemical Additive Manufactured Heat Sink Designs for Single-Phase Immersion Cooling

by

Joseph Herring

DISSERTATION

Submitted in partial fulfillment of the requirements
for the degree of Doctor of Philosophy at
The University of Texas at Arlington
August 2024

Arlington, Texas

Supervising Committee:

Dereje Agonafer, Supervising Professor
Abdolhossein Haji-Sheikh
Miguel Amaya
Amir Ameri
A S M Raufur Chowdhury

Acknowledgements

I am deeply grateful to Dr. Dereje Agonafer for his unwavering support, invaluable guidance, and profound belief in my potential. Your investment in my academic and personal growth has been a cornerstone of this journey.

I extend my sincere thanks to my Doctoral Committee for their dedication, time, and insightful feedback, which have been instrumental in shaping this work.

I also wish to acknowledge my academic and professional colleagues. It has been an honor to work alongside individuals whom I not only respect professionally but am fortunate to call great friends. Your camaraderie and shared wisdom have enriched this experience immeasurably.

To my Mother and Father, your constant belief in my abilities and encouragement to push beyond my perceived limits have been a source of strength throughout my life.

Lastly, my deepest appreciation goes to my wife, Katherine. Your unwavering support and love, for me and our four children, have made the completion of this milestone possible.

August 12, 2024

Dedication

To my Katherine, and Markiah Grace, Charlotte Rae, Clara Elizabeth, and Joseph Frank III.
Know that it is all for you.

Abstract

Traditional air-cooling along with corresponding heat sinks are beginning to reach performance limits, requiring lower air-supply temperatures and higher air-supply flowrates, in order to meet the rising thermal management requirements of high power-density electronics. A switch from air-cooling to single-phase immersion cooling provides significant thermal performance improvement and reliability benefits. When hardware which is designed for air cooling is implemented within a single-phase immersion cooling regime, optimization of the heat sinks provides additional thermal performance improvements. This work investigates performance of a machine learning (ML) approach to building a predictive model of the multi objective and multi-design variable optimization of an air-cooled heat sink for single-phase immersion-cooled servers. Parametric simulations via high fidelity CFD numerical simulations are conducted by considering the following design variables composed of both geometric and material properties for both forced and natural convection: fin height, fin thickness, number of fins, and thermal conductivity of the heat sink. Generating a databank of 864 points through CFD numerical optimization simulations, the data set is used to train and evaluate the machine PhD Dissertation Defense Announcement Mechanical and Aerospace Engineering Department University of Texas at Arlington learning algorithms' ability to predict heat sink thermal resistance and pressure drop across the heat sink. Three machine learning regression models are studied to evaluate and compare the performance of polynomial regression, random forest, and neural network to accurately predict heat sink thermal resistance and pressure drop as a function of various design inputs. This approach to utilizing numerical simulations for building a databank for machine learning predictive models can be extrapolated to thermal performance prediction and parameter optimization in other electronic thermal management applications and thus reducing the design

lead time significantly. Heat sinks designed for electrochemical additive manufacturing (ECAM) with Body Centered Cubic (BCC) lattice structures are evaluated using computational fluid dynamics (CFD) conjugate heat transfer (CHT) analyses in ANSYS Fluent for single-phase immersion cooling applications. More complex heat sink cooling surface geometries enabled by ECAM fabrication technologies have a greater surface area to volume ratio than traditional parallel plate fins. To benchmark performance, we establish a baseline immersion cooling heat sink metric for various dielectric fluid flowrates using a conventional finned heat sink. We then compare the thermal resistance and pressure drop characteristics of this baseline with those of the ECAM BCC lattice heat sink design. Additional design factors of wall thickness and porosity are also considered. This study evaluates the thermal performance of ECAM-fabricated BCC lattice heat sinks as an innovative solution for enhancing cooling efficiency in high power-density electronics immersion cooling applications. The findings are expected to offer valuable insights into the viability and performance advantages of such heat sinks. By leveraging the capabilities of AM designed structures, this research contributes to the development of more effective and sustainable immersion cooling solutions for next generation electronic systems.

Joseph Herring, Ph.D.

The University of Texas at Arlington, 2024

Supervising Professor: Dereje Agonafer

List of Illustrations

List of Tables

Nomenclature

β	Volumetric Expansion Coefficient
μ	Dynamic Viscosity
ρ	Density
ρ_0	Constant Density
CoP	Coefficient of Prognosis
$\vec{F}$	External Force
$\vec{g}$	Gravity
h	Sensible Enthalpy
k	Molecular Conductivity
k_t	Turbulence Transport Conductivity
S_h	Volumetric Heat Generation
SS_E	Regression Variation
SS_T	Total Variation
T	Temperature
T_0	Operating Temperature
t	Time
$\vec{v}$	Velocity
$\mathbf{x_i}$	Model Data Input
$\boldsymbol{y_i}$	Data Set Target Output
$\hat{y}$	Model Predicted Output

Table of Contents

Chapter 1

Introduction

1.1 Background and Motivation

1.1.1 The Need for Data Center Energy Efficiency

The rapid growth of digital technology and the increasing reliance on data-intensive applications have driven a significant rise in data center energy consumption. The International Energy Agency (IEA) recently projected that global data center electricity demand will more than double by 2026 [1]. This surge in energy use is attributed to the proliferation of cloud computing, artificial intelligence, and other high-performance computing applications, which require substantial processing power and, consequently, more energy to operate. As data centers become larger and more numerous, their energy consumption is becoming a critical concern for both operators and policymakers, necessitating urgent action to enhance energy efficiency.

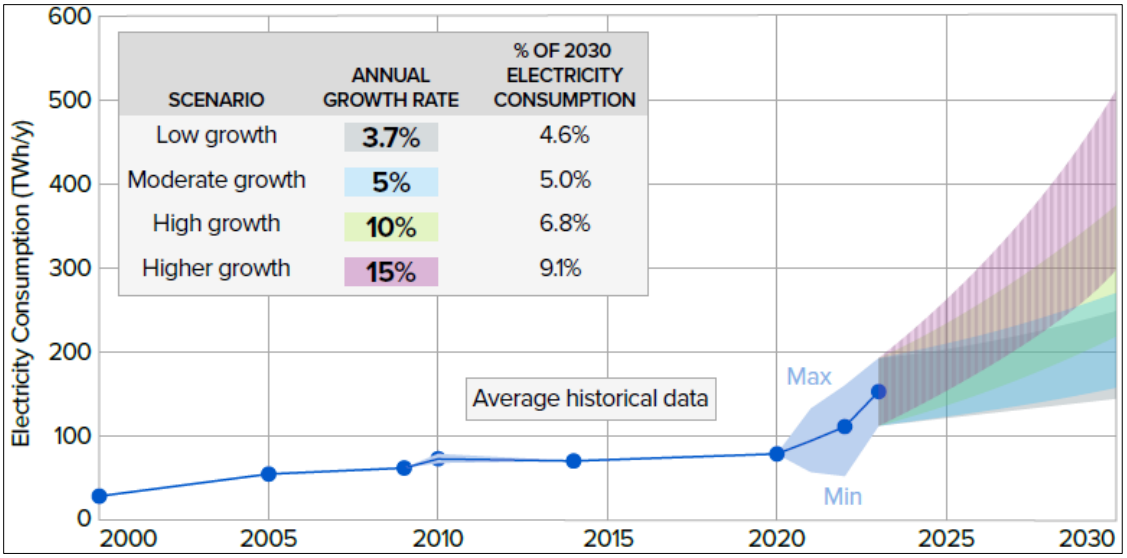

Figure 1.1 US Data Center 2030 Power Consumption Projections [1]

While predictions can be made about the future growth potential of the data center industry and its energy consumption, the emerging proliferation of cloud-based Artificial Intelligence applications both in industrial business use and consumer daily life poses an unseen potential for exponential growth and future energy demand requirements. Although current estimates show that AI applications use only 10-20% of data center electricity, this number is readily increasing.

AI presents an additional challenge in that it is significantly more energy-intensive than traditional data center applications which have drove the industries growth over the last twenty years. Recent data suggests that a single ChatGPT request at 2.9 watt-hours requires nearly 10 times more energy than a typical 0.3 watt-hour Google query. There is no precedent set to know how future energy demand will increase with the rise of generative AI text, images, audio, and other media. One of the primary challenges contributing to the escalating energy demand is the continuous advancement in computing and processing power. As chips become smaller and more powerful, they generate significantly higher heat fluxes, which in turn increases the thermal load on data center cooling systems. Traditional air-cooling methods are becoming less effective at managing these higher heat outputs, leading to inefficiencies and higher energy consumption. To address this issue, innovative thermal management solutions are required to maintain the performance and reliability of data center operations while minimizing energy use.

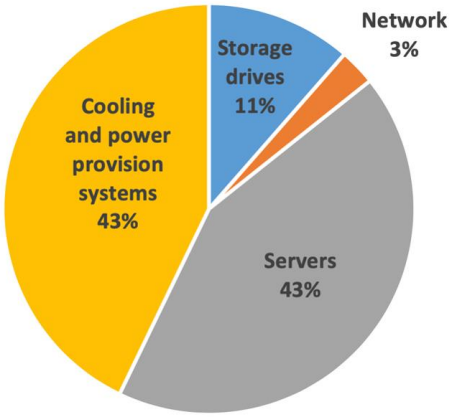

Figure 1.2 Fractional Breakdown of Data Center Energy Consumption [2]

Single-phase immersion cooling has emerged as a promising energy-efficient and reliable thermal management solution to address the challenges of increasing component heat dissipation. By submerging electronic components in a dielectric fluid, this method allows for more effective heat transfer and cooling compared to traditional air or water-cooling techniques. The dielectric fluid efficiently absorbs and dissipates heat, reducing the need for energy-intensive fans and air conditioning units. Furthermore, single-phase immersion cooling offers the added benefit of reducing the physical footprint of cooling infrastructure, enabling more compact and efficient data center designs.

As data centers continue to expand in both size and capacity, the need for energy-efficient cooling solutions becomes increasingly urgent. Implementing single-phase immersion cooling not only helps to manage the higher heat fluxes generated by modern processors but also contributes to reducing the overall energy consumption of data centers. By adopting such advanced cooling technologies, data center operators can meet the growing demand for processing power while

simultaneously addressing the critical need for energy efficiency, ultimately supporting global sustainability goals and reducing the environmental impact of digital infrastructure.

1.1.2 The Need for High-Performance Electronics Thermal Management

For nearly six decades, the semiconductor industry has adhered to Moore's Law, a guiding principle that states, "The number of transistors in an integrated circuit will double every 18 months." This trend has driven remarkable advancements in computing and processing power, allowing for the development of increasingly sophisticated and compact electronic devices. However, as transistor densities continue to rise, so too does the power density within semiconductor packages. This escalation in power density leads to greater component heat fluxes, presenting significant challenges in maintaining the thermal stability and performance of high-performance electronics.

The exponential increase in heat generated by modern semiconductor devices has made thermal management a critical factor in the design and operation of electronic packages. As processing power increases, so does the energy consumption and heat output of these devices. Without effective thermal management, the excessive heat can lead to reduced performance, shortened component lifespans, and, in extreme cases, catastrophic failure. Thus, package thermal management has become a key limiting factor for semiconductor manufacturers, directly

impacting the efficiency, reliability, and longevity of central processing units (CPUs) and other critical components.

To address these challenges, the industry must prioritize the development and implementation of advanced thermal management solutions. Traditional cooling methods, such as air or basic liquid cooling, are often inadequate to handle the high heat fluxes generated by modern semiconductor devices. This inadequacy necessitates the adoption of more sophisticated techniques, such as single-phase immersion cooling, advanced heatsink designs, and thermoelectric cooling systems. These methods are designed to efficiently dissipate heat from high-density packages, ensuring that the components operate within safe temperature ranges and maintain their performance over time.

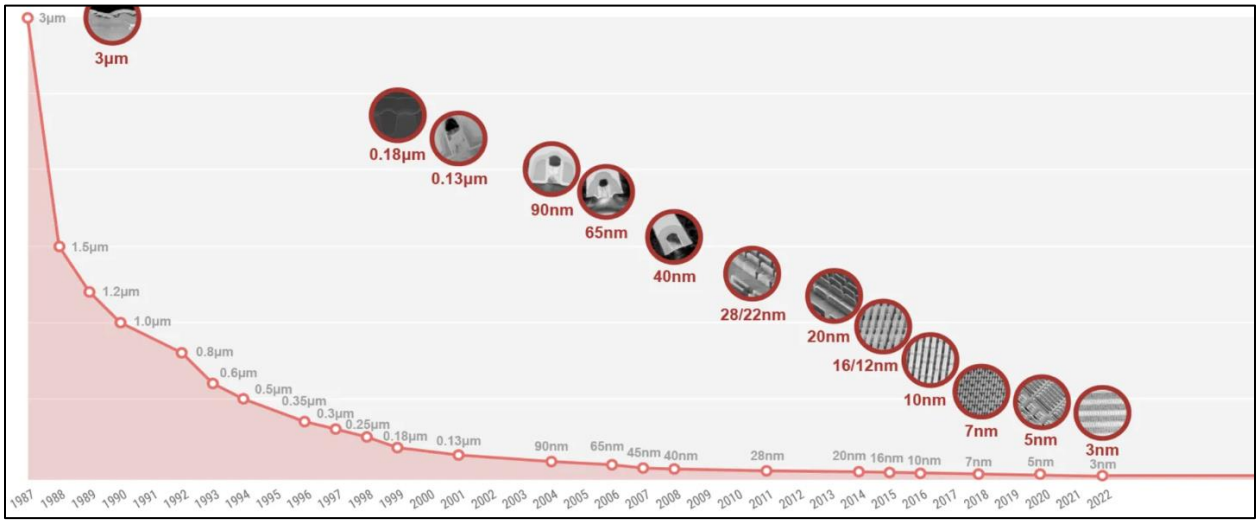

Figure 1.3 Transistor Size over Time [3]

The need for high-performance electronics thermal management is more pressing than ever as the industry continues to push the boundaries of what is possible with semiconductor technology. By investing in innovative thermal management solutions, manufacturers can overcome the thermal limitations imposed by Moore's Law and continue to deliver the powerful, efficient, and reliable electronic devices that drive today's technological advancements. In doing so, they will not only meet the demands of current applications but also pave the way for future innovations in computing and electronics.

1.2 Data Center Introduction

Data centers are specialized facilities designed to house the essential computing, storage, networking, power, and cooling infrastructure required to operate computer servers effectively. These servers, which include both compute and storage hardware, are the backbone of modern digital infrastructure, enabling the processing, storage, and transmission of vast amounts of data. As such, data centers play a critical role in supporting a wide range of applications, from large-scale data storage to cloud computing and internet connectivity. They are the critical infrastructure of the modern world that is required to power everything from online services and social media platforms to enterprise-level business applications and scientific research.

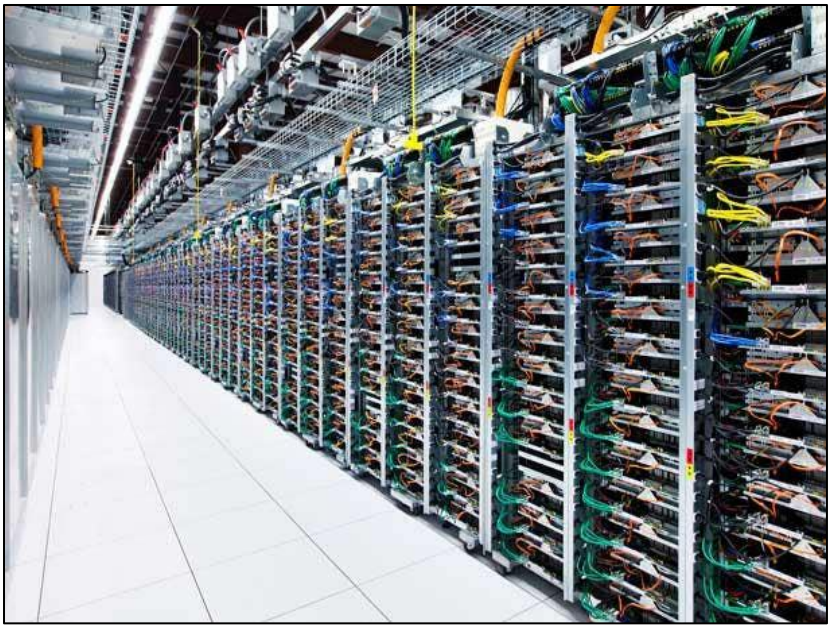

Figure 1.4 Data Center Server Rack Aisle [4]

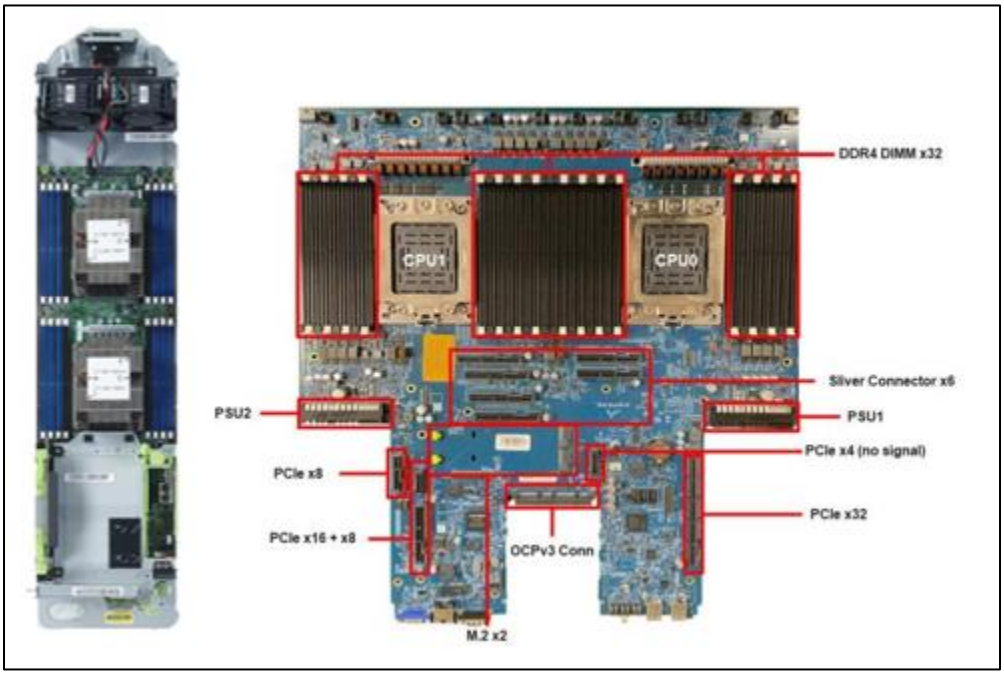

Figure 1.5 Shadow Core vs. Spread Core Server Design (E7278 & Mt. Jade Ampere Altra) [5-6]

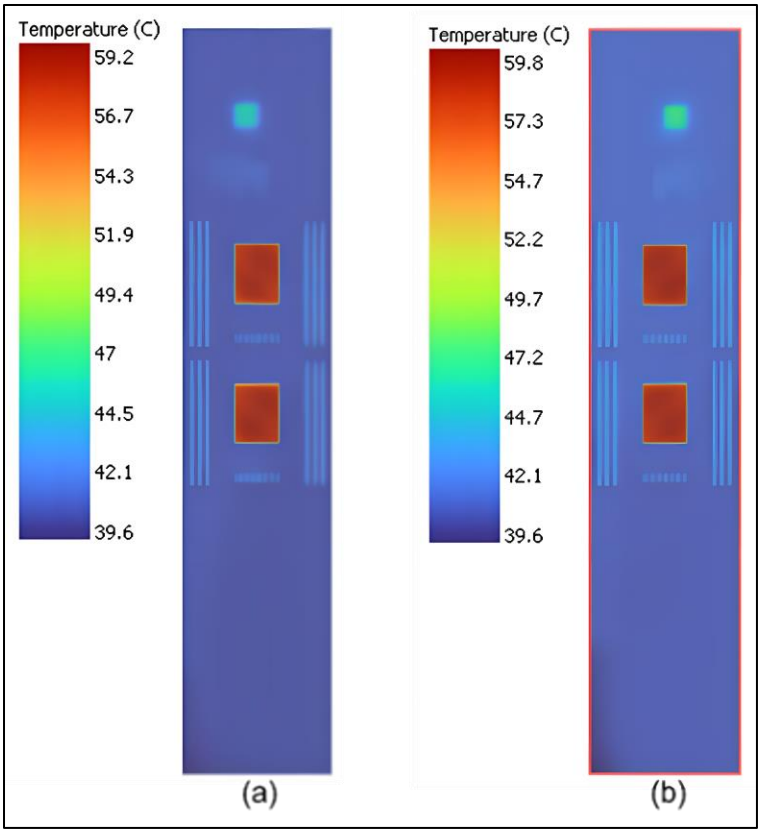

Figure 1.6 Thermal Profiles for (a) tank configuration and (b) sled configuration [7]

1.3 Single-phase Immersion Cooling

Single-phase immersion cooling involves submerging electronic components, such as servers and processors, directly into a dielectric fluid. This fluid is non-conductive, allowing it to come into direct contact with the electronic components without causing short circuits or electrical signal integrity issues. The cooling fluid possesses a greater specific heat capacity than that of air, allowing for increased thermal management performance. This method of cooling offers a more efficient and uniform heat transfer compared to traditional air or liquid cooling systems, which often struggle to keep up with the high heat fluxes generated by modern high-performance electronics.

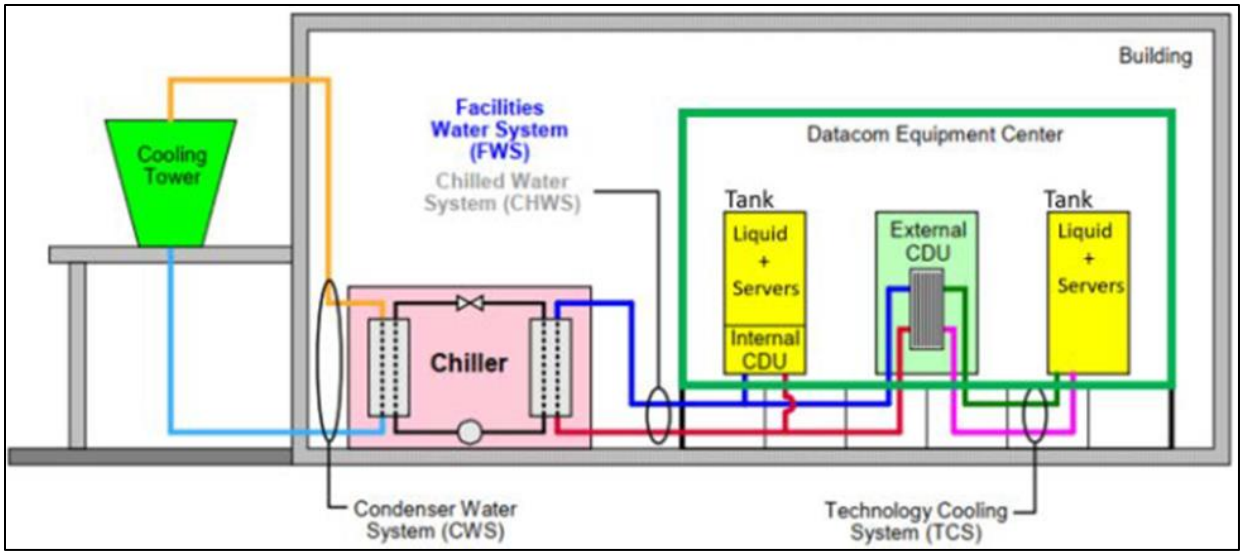

Figure 1.7 Single-phase Immersion Cooling Diagram [8]

The adoption of single-phase immersion cooling has been driven by the need to manage the increasing thermal loads in data centers and other computing environments. As processors and IT compute equipment such as CPUs and GPUs become more power-dense, traditional cooling methods become less effective, leading to potential overheating and reduced performance. Single-phase immersion cooling addresses these challenges by providing a highly efficient, reliable, and scalable solution. Not only does it enhance the thermal performance of the systems, but it also reduces energy consumption by eliminating the need for energy-intensive cooling fans and air conditioning units. This makes it an attractive option for data centers looking to improve their energy efficiency and reduce operational costs while maintaining the performance and critical infrastructure reliability.

Chapter 2

Computational Fluid Dynamics (CFD) Evaluation of Electrochemical Additively Manufactured Heat Sinks for Single-Phase Immersion Cooling

Reprinted with Permission [30]

2.1 Introduction

Advances in computing and processing power along with smaller chip sizes continue to generate greater component heat fluxes [1]. Effective thermal management is critical to maintaining device performance, reliability, and safety [2]. Standard practices involving traditional electronics cooling methodologies, such as forced-air cooling, are approaching a threshold and will not continue to be suffice the escalating power-densities of high-performance electronics [3]. Energy costs are approximately 50% of a typical data center's total operating cost, and in general consume 100 – 200 times more energy than an office building [4]. This necessitates the development and implementation of advanced cooling solutions to efficiently meet thermal performance requirements. High performance computing, artificial intelligence, computing at the edge, and modern infrastructure all rely on data centers for the storage and processing of data.

2.1.1 Background and Motivation

Improvements in air cooling technologies have allowed data centers to advance computing and storage capabilities [5-7]. However fundamental limitations still exist at the expense of reliability, efficiency, sustainability, space utilization, increased noise, and significant energy

consumption and costs [8-10]. Air-cooled systems are susceptible to environmental factors such as moisture and air contaminants, which contribute to mechanical failures [11-13]. Liquid cooling technologies provide more heat dissipation, more space and energy efficiency, and increased reliability than air-cooling [7,14-16].

Two of the primary emergent liquid cooling technologies for data center cooling applications are immersion cooling and direct-to-chip cooling [9]. Single-phase immersion cooling (SPIC) places the server electronics directly in contact with a dielectric cooling fluid with a much greater thermal capacity than air. SPIC also improves energy efficiency and system reliability by preventing airborne contaminants and lowers OpEx by eliminating air-cooling fans and air handling units.

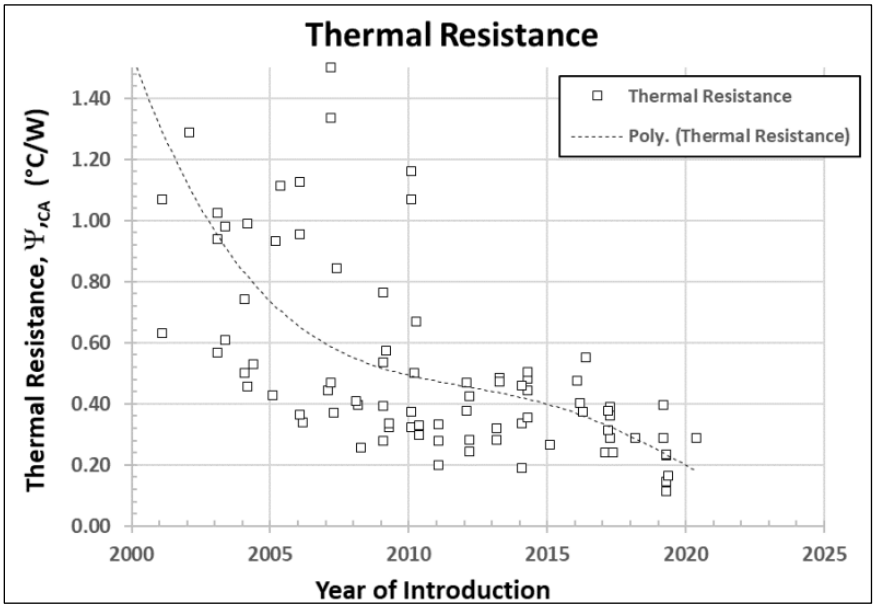

Figure 2.1 ASHRAE Thermal Resistance Required to Cool Socket Power [9]

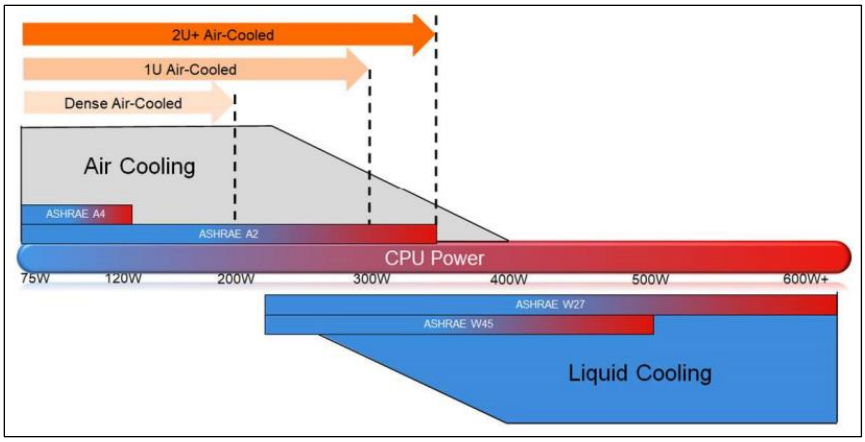

Figure 2.2 ASHRAE TDP Cooling Performance Comparison [9]

Unlike direct-to-chip liquid cooling, SPIC does not require complex liquid distribution manifold systems. There are still design challenges that must be addressed for a SPIC system. The adoption of natural convection or forced flow immersion cooling in existing data centers requires optimizing and retrofitting air-cooled server components to take full advantage of the potential performance improvements of immersion cooling [17].

2.1.2 Objectives

The primary goal of this research study is to assess the thermal performance of heat sink designs enabled by electrochemical additive manufacturing (ECAM) for application in single-phase immersion cooling systems.

This study aims to derive heat sink design parameters governing the thermal performance such as porosity and wall thickness and identify boundary conditions that simulate real-world

application to derive insight into the performance improvements brought by additively manufactured heat sink geometries.

This study establishes a baseline for thermal performance using a traditional finned heat sink design to serve as a reference point for accurate comparison and assessment of heat sink performance.

2.1.3 Literature Review

An increase has been shown in single-phase immersion cooling thermal performance with the use of heat sinks designed and optimized for dielectric fluid and flow rates utilized in natural convection and forced flow immersion cooling regimes [18-21]. Advances in additive manufacturing (AM) technologies have enabled heat sink geometries that previously could not be manufactured by traditional techniques. Heat transfer improvements enabled by AM are increasingly being explored within academic literature. These cooling surface geometries made capable by AM include Triply Periodic Minimal Surfaces (TPMS) [22,23], lattice structures [24,25], and topology optimized surfaces [26,27].

2.2 Heat Sink Design Overview

The heat sinks evaluated in this study are designed to increase the heat sink cooling surface area with ECAM manufacturability in mind. The ECAM process is a metal additive manufacturing process based on electrodeposition at the atomic scale, in a similar fashion to common copper

electroplating techniques. A localized electric field is generated in an electrolyte feedstock by pixels which activate to selectively deposit metal and builds up in layers [28].

ECAM enables parts to be manufactured with much higher resolution and lower surface roughness than other common metal additive manufacturing processes, with a minimum feature size of 70 μm. Other standard metal additive post processing requirements such as curing, sintering, and de-powdering are not required, reducing part print times and costs.

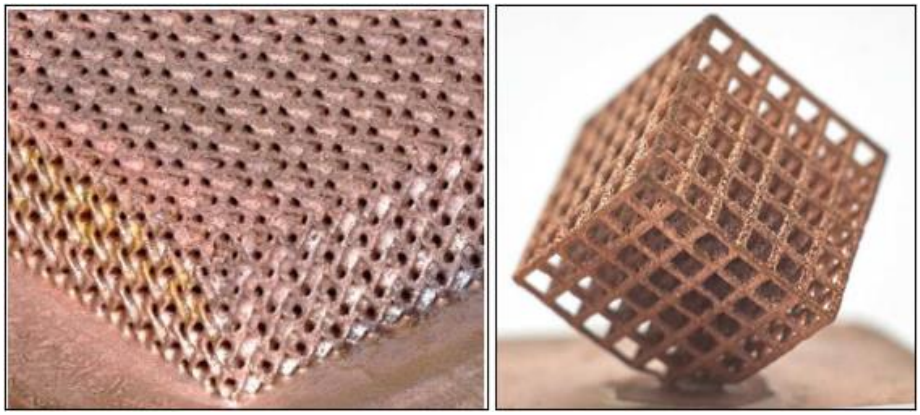

Figure 2.3 ECAM Printed Copper TPMS and Lattice Geometries [28]

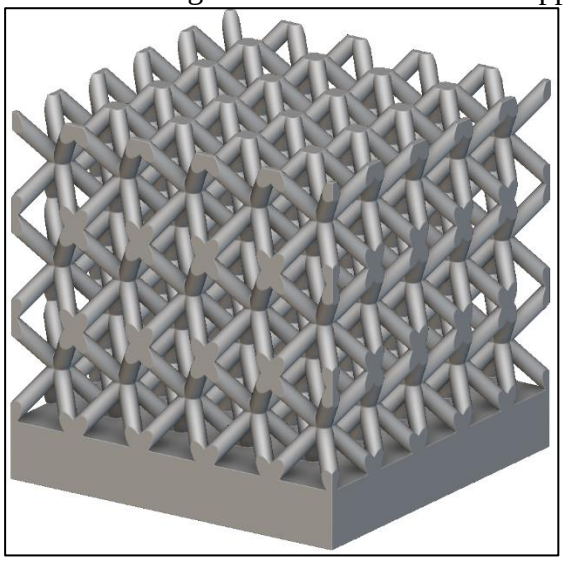

Figure 2.5 Body Centered Cubic Lattice Heat Sink Design

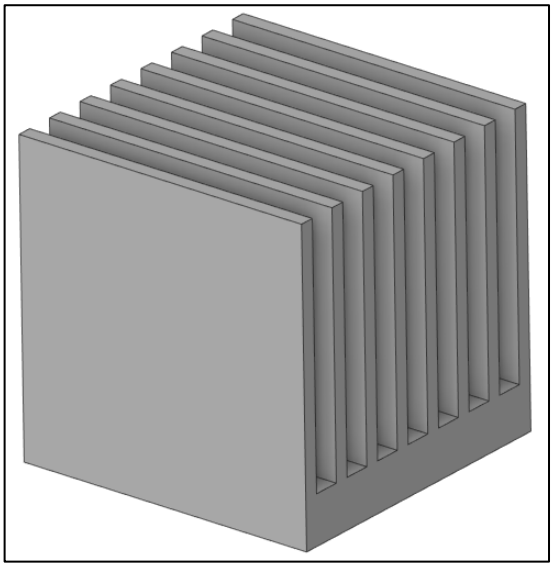

Figure 2.4 Baseline Channel Fin Heat Sink Design

Table 2.1 Heat Sink Design Constraints

Parameter	Value
Base Height	4 mm
Base Width	25 mm
Base Length	25 mm
Lattice Height	21 mm
Fin/Lattice Wall Thickness	1.2 mm
Component Heater Area	15mm x 15mm

Table 2.2 Heat Sink Design Variables

Geometry	Porosity, λ
Traditional Fins	60% (2.20 mm fin gap)
	70% (3.56 mm fin gap)
	80% (6.73 mm fin gap)
BCC Lattice	60%
	70%
	80%

Table 2.3 BCC Lattice Unit Cell Sizes

Porosity, λ	Unit Cell Size
60%	3.6mm x 3.6mm x 3.9mm
70%	4.2mm x 4.2mm x 4.8mm
80%	5.1mm x 5.1mm x 7.0mm

In this study, a Body Centered Cubic (BCC) lattice heat sink was designed and analyzed. The heat sink dimensions are 25 mm x 25 mm, with a 4 mm thick base, 21 mm lattice fin height and 1.2 mm wall thickness. Heat sink designs were generated with lattice porosities of 60%, 70%, and 80% to evaluate the impact of porosity on thermal performance and pressure drop. Traditional channel fin heatsinks were also created with equivalent porosities and fin wall thickness. ANSYS SpaceClaim was used to create the channel fin heat sink geometries. The BCC lattice heat sink

designs were created using an implicit modeling approach with nTopology. Implicit geometries for the lattice heat sink designs were exported as STL files.

The BCC Lattice heat sink model is shown in Fig. 2.4 and the Baseline Channel Fin heat sink model is shown in Fig. 2.5. Table 2.1 contains the fixed design constraints for both heat sink types and Tables 2.2 and 2.3 contain the heat sink design variables for 60%, 70%, and 80% porosity design permutations.

2.3 Computational Fluid Dynamics Methodology

ANSYS Fluent 2023 R2 was used to create a computational fluid dynamics model fluid domain, mesh the heat sink solid models and fluid domain, and solve the conjugate heat transfer and flow equations. The solid domain for both the BCC lattice heat sinks and fin baseline heat sinks were given material properties of pure copper, and the fluid domain was given material properties of PAO-6 dielectric fluid.

2.3.1 CFD Model Setup and Boundary Conditions

The CFD model setup and boundary conditions are shown in Fig. 2.6 and tabulated in Table 2.4.

Table 2.4 CFD Model Boundary Conditions

Parameter	Boundary Type	Value(s)
Solid Domain	Material	Copper
Fluid Domain	Material	PAO-6
Walls	Material	Acrylic
Inlet	Mass Flow Inlet	0.045 kg/s @ 40 °C
Outlet	Pressure Outlet	-
Heat Load	Wall	150 W

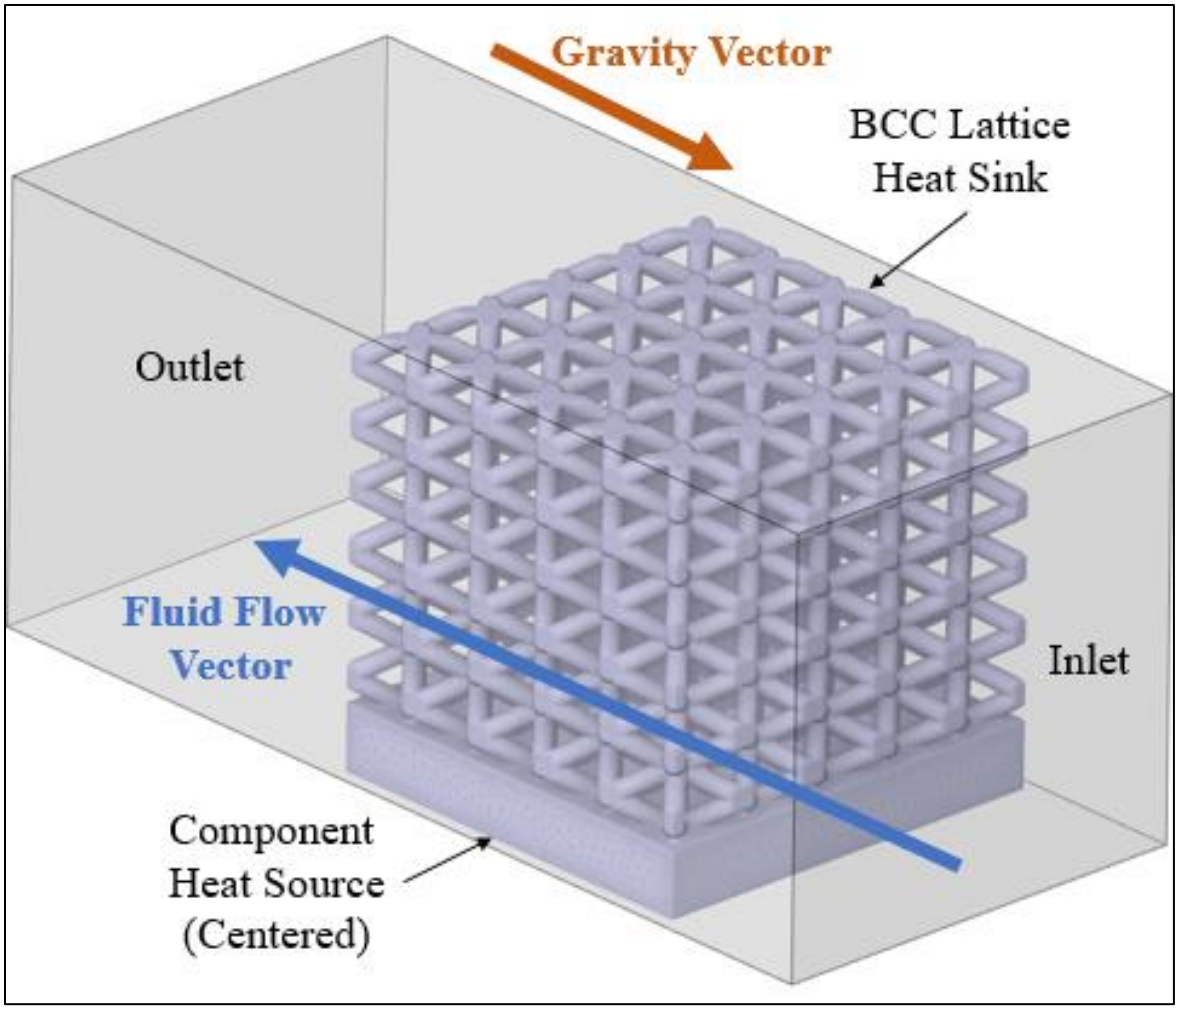

Figure 2.6 Heat Sink CFD Model Setup and Boundary Conditions

The fluid domain size and dielectric fluid inlet flow rate boundary conditions were modeled after a real-world forced convection heat sink design for single-phase immersion cooling that utilizes inlet propellers and an enclosure around the heat sink to significantly increase dielectric fluid flow rates directly at the heat sink inlet and contain the fast-moving fluid as it passes through the cooling surfaces of the heat sink [29].

2.3.2 CFD Model Mesh and Validation

The CFD model mesh was generated with the ANSYS Fluent 2023 R2 meshing. For the BCC lattice heat sink models with faceted STL geometry, the CFD model was specified to contain faceted bodies within the water-tight meshing tool, creating a conformal mesh between the solid and fluid domains.

The volume mesh was generated using polyhedral mesh elements, and a cross-sectional view of the 70% porous BCC lattice heat sink model mesh is shown in Fig. 2.7. Table 2.5 contains the mesh element count for all CFD model meshes.

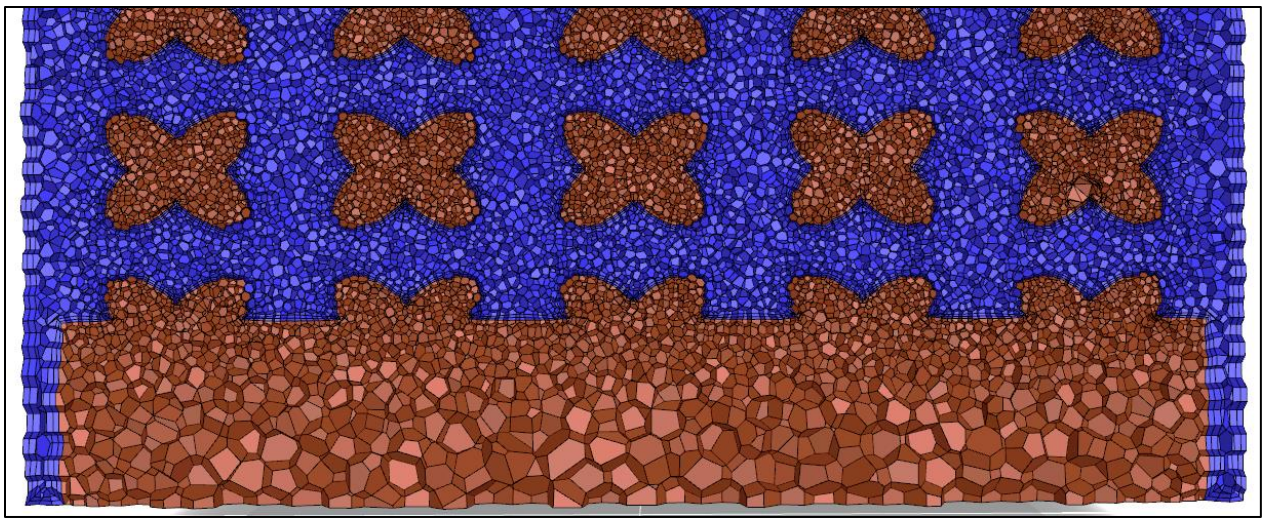

Figure 2.7 70% Porosity BCC Lattice Heat Sink CFD Volume Mesh

A mesh independence study was conducted to validate convergence of the CFD model. The mesh convergence was evaluated for the 70% porosity finned heat sink. Fig. 2.8 compares the junction temperature with increasing heat sink volume mesh element count. The volume mesh resulting in 346,034 elements was used to establish surface and volume meshing parameters for all finned heat sink CFD model meshes.

Table 2.5 CFD Model Volume Mesh Element Count

Heat Sink	Porosity, λ	Mesh Element Count
Traditional Fins	60%	295,358
	70%	346,034
	80%	78,028
BCC Lattice	60%	9,477,520
	70%	2,858,681
	80%	2,055,655

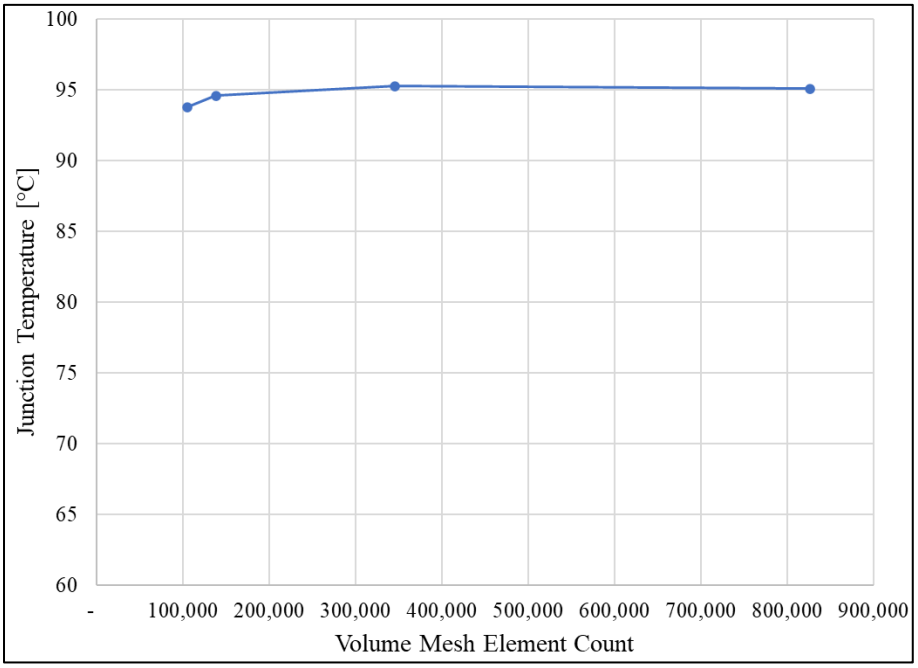

Figure 2.8 CFD Model Mesh Independence Study

2.3.3 CFD Model Conditions and Equations

The CFD model included the energy equation and viscous effects were considered using the SST k-ω model. Gravity effects are considered as shown in Fig. 2.6 to represent a data center server SPIC setup orientation. The Coupled solution scheme solver was used with all standard under-relaxation factors. The governing equations for the flow are listed below in (2.1) through (2.5).

Conservation of Mass:

$$\frac{\partial \rho}{\partial t} + \nabla \cdot (\rho \vec{v}) = 0 \tag{2.1}$$

For incompressible flow, (2.1) reduces to:

$$\nabla \cdot (\vec{v}) = 0 \tag{2.2}$$

Conservation of Momentum:

$$\rho \frac{\partial \vec{v}}{\partial t} + \rho(\vec{v} \cdot \nabla)\vec{v} = -\nabla p + \rho \vec{g} + \mu \nabla^2 \vec{v} \tag{2.3}$$

Conservation of Energy for the fluid domain:

$$\rho \frac{\partial h}{\partial t} + \rho \nabla \cdot (h\vec{v}) = -\nabla p + \nabla \cdot [(k + k_t)\nabla T] + S_h \tag{2.4}$$

Solid domain energy equation:

$$\frac{\partial}{\partial t}(\rho h) = \nabla \cdot (k \nabla T) + S_h \tag{2.5}$$

2.4 Results and Discussion

Conjugate heat transfer and fluid flow computational fluid dynamics analysis was completed for the heat sink designs tabulated in Table 2.2 using the methodology outlined in Section III. The following results compare the thermal performance of the BCC lattice heat sinks designed for electrochemical additive manufacturing to the established channel fin baseline heatsinks.

The resulting maximum junction temperatures are shown in Fig. 2.15, and the resulting pressure drop across the heatsink inlet and outlet are shown in Fig. 2.16. Heat sink temperature contour plots are shown from Fig. 2.9 – 2.14. Table 2.6 tabulates the heat sink thermal performance increase for each heat sink porosity as a percentage of maximum junction temperature reduction from the baseline finned heat sink.

For all heat sink porosities considered, the BCC lattice heatsink reduced maximum junction temperature compared to the fin baseline heat sinks.

The BCC lattice heat sinks reduced the maximum junction temperature by 50.5% at 80% porosity, 21.2% at 70% porosity, and 13.4% at 60% porosity when compared to the respective channel fin baseline heat sinks.

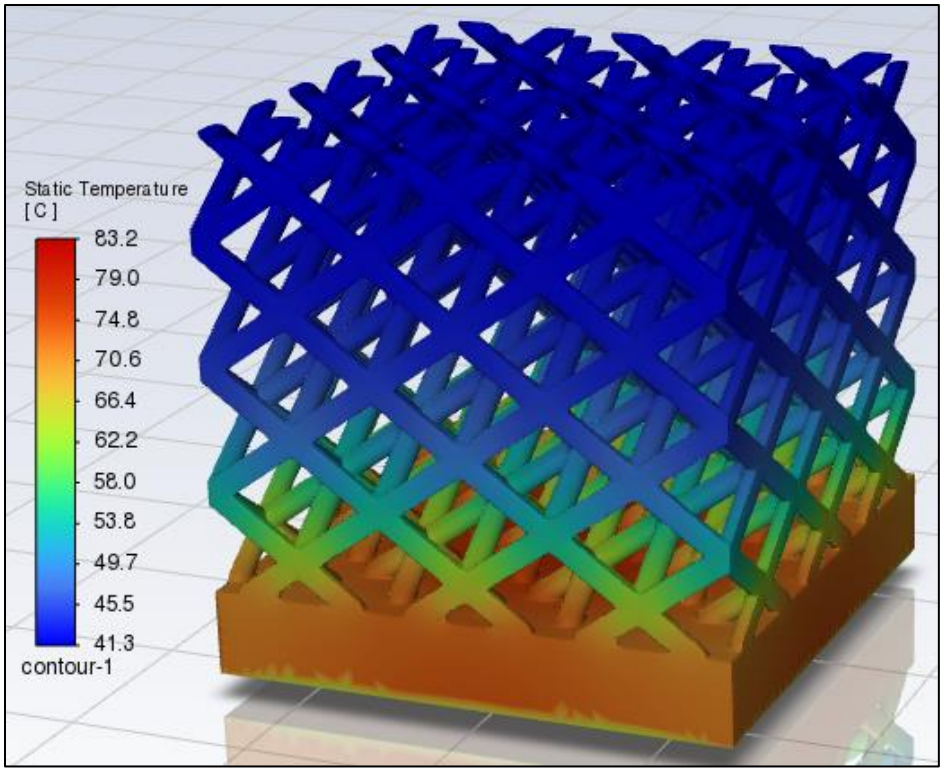

Figure 2.9 80% Porosity BCC Lattice Heat Sink

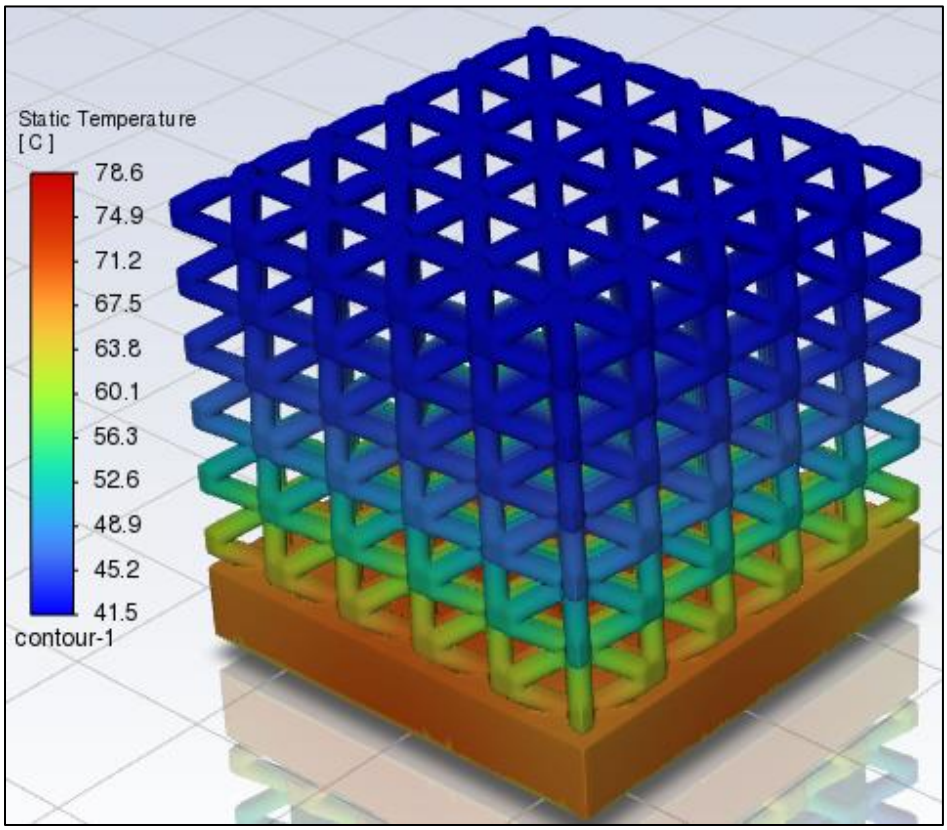

Figure 2.10 70% Porosity BCC Lattice Heat Sink

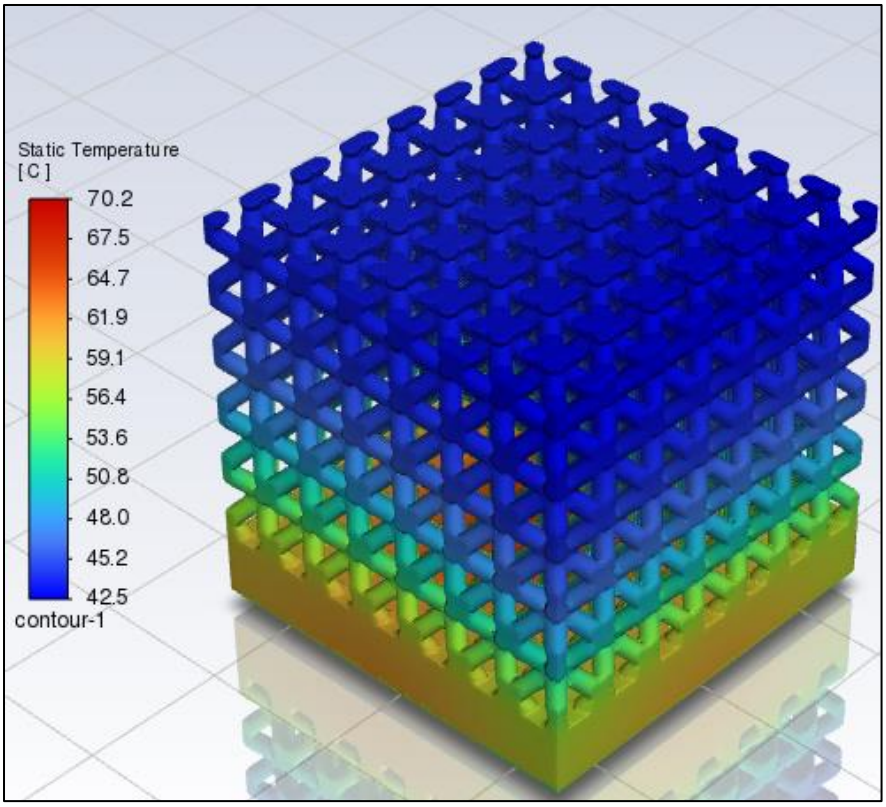

Figure 2.11 60% Porosity BCC Lattice Heat Sink

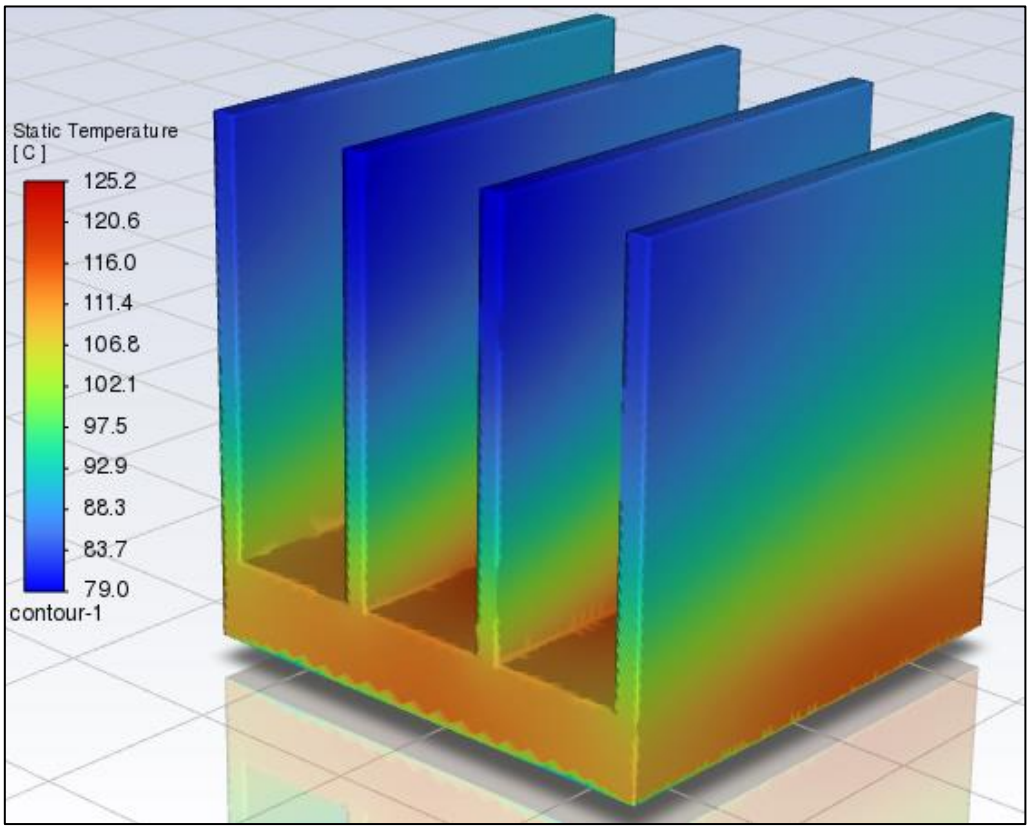

Figure 2.12 80% Porosity Baseline Fin Heat Sink

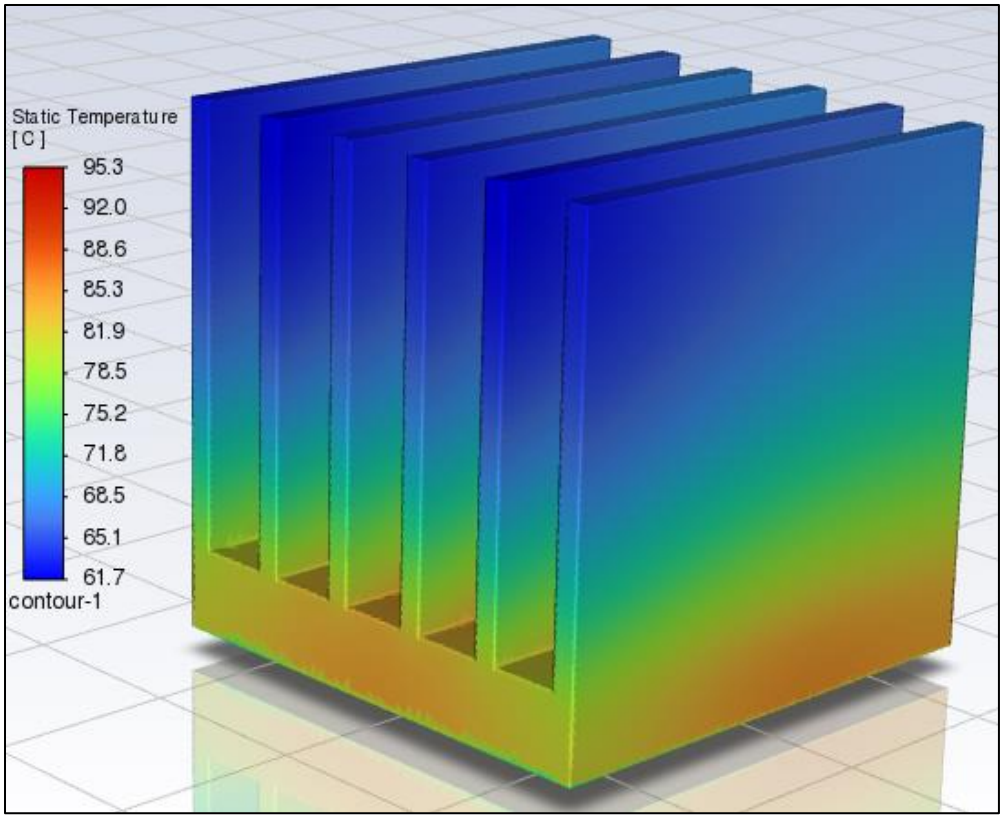

Figure 2.13 70% Porosity Baseline Fin Heat Sink

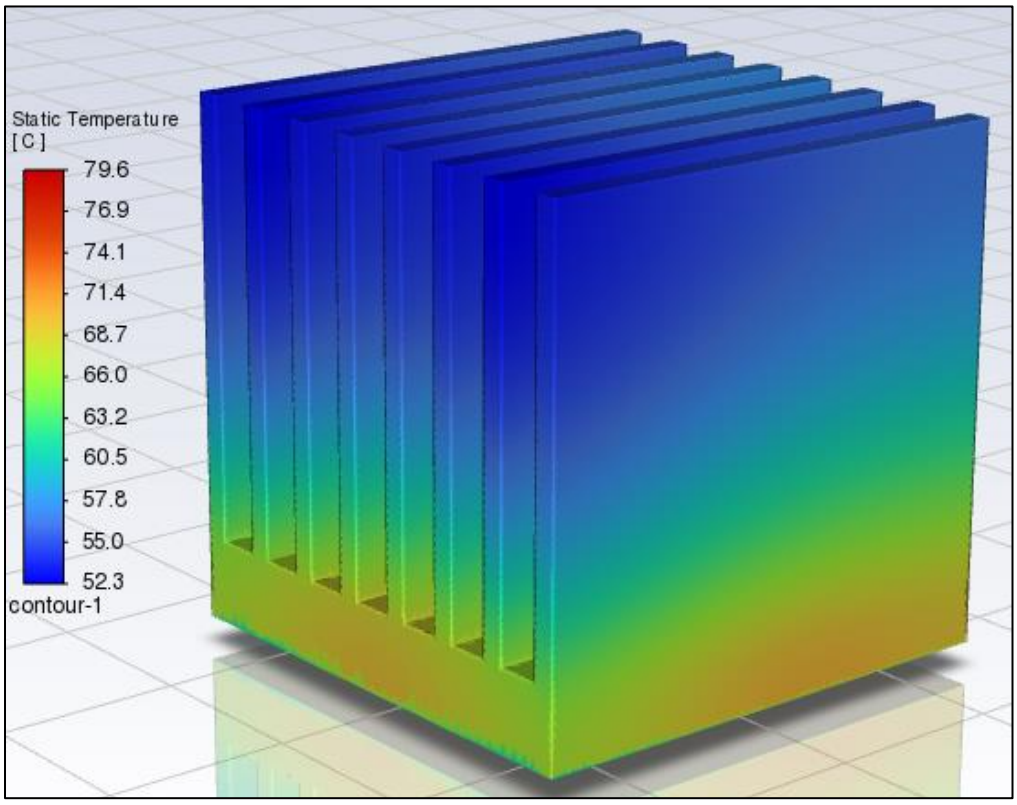

Figure 2.14 60% Porosity Baseline Fin Heat Sink

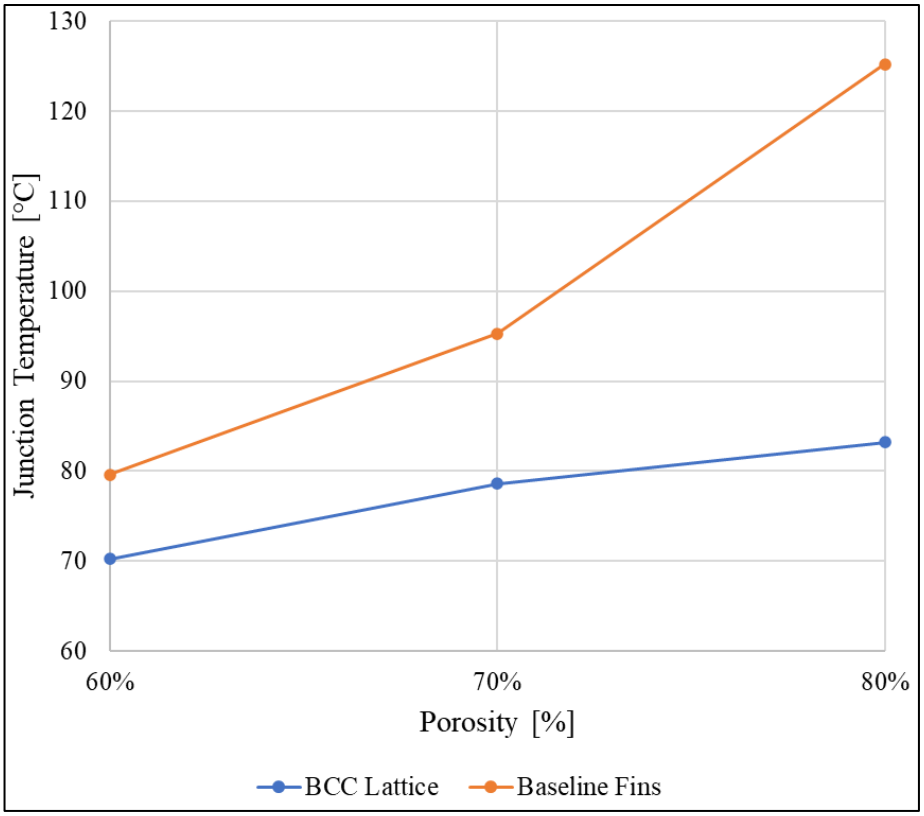

Figure 2.15 Junction Temperature of BCC Lattice vs. Baseline Heat Sinks

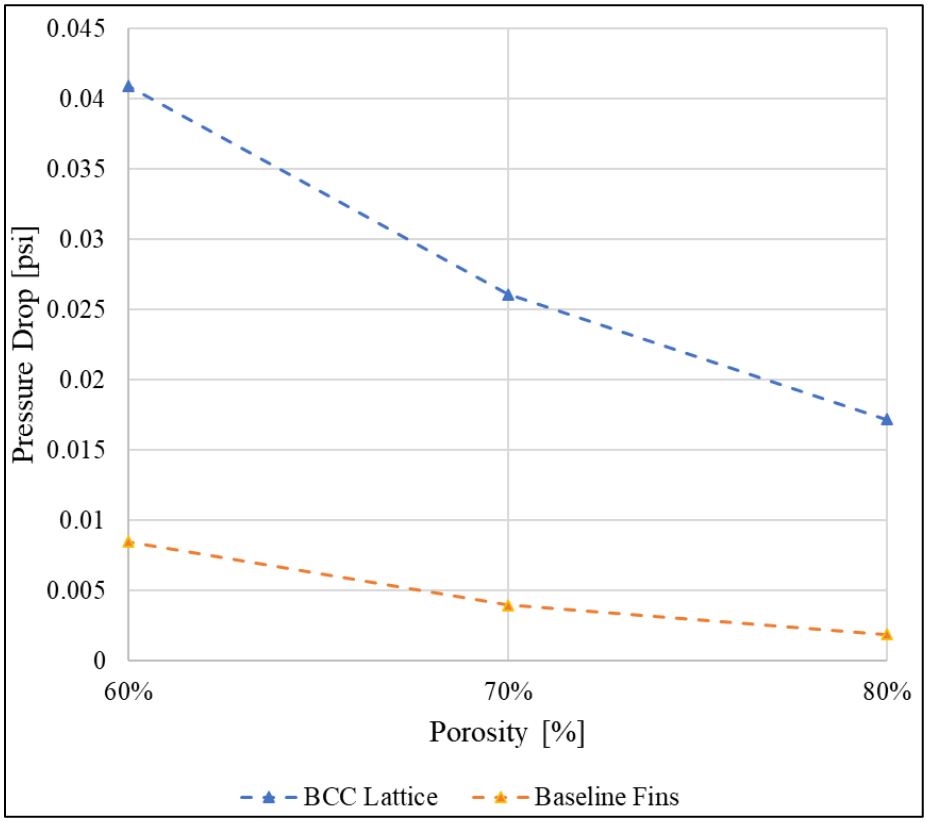

Figure 2.16 Pressure Drop of BCC Lattice vs. Baseline Heat Sinks

Table 2.6 Junction Temperature Reduction BCC vs Baseline

Porosity, λ	Max Junction Temperature Reduction (°C)
60%	9.4
70%	16.7
80%	42

2.5 Conclusion

Heat sinks with improved cooling surfaces are made possible with advances in additive manufacturing technologies such as ECAM. BCC lattice structure geometries significantly improved the thermal performance of a heat sink designed for forced convection single-phase immersion cooling. The BCC lattice heat sinks enhanced thermal performance with a reduction in component junction temperature for all porosities considered.

ECAM enabled geometries can be utilized in other potential liquid and air-cooling regimes, with full application-specific customization and more untapped optimization potential. Additionally, other cooling surfaces may be explored as potential heat sink designs now possible is expanding.

To expand on this work, further investigation into additional porosities, wall thicknesses, lattice types, and TPMS heatsinks for electronics cooling applications would provide further insight into the potential thermal performance benefits that ECAM and complex heat sink designs offer. Future experimental validation of the results herein, as well as other geometries and design parameters will also provide further advancement of this technology.

2.6 Acknowledgements

We would like to acknowledge Fabric8Labs for sponsoring this work, and Yeong-Yan Perng of the ANSYS Fluent team for his contributions to creating a functional Fluent workflow for these heat sink designs.

Chapter 3

Single-Phase Immersion Cooling Multi-Design Variable Heat Sink Optimization for Natural Convection

Reprinted with Permission [13]

3.1 Introduction

Liquid immersion cooling technology has proven the potential to provide improved thermal management for data center computation and storage systems [1]. Fully submerging server systems in dielectric fluid provides improved heat transfer, reduced operating noise, and increased energy efficiency over forced-air cooling [2]. The adoption of immersion cooling in existing data centers requires additional retrofitting of server components to take full advantage of the potential performance improvements of immersion cooling [3]. Heat sinks are critical in effectively dissipating heat from electronic devices to the surrounding dielectric fluid in an immersion cooling system.

An optimal heat sink fin configuration enhances heat transfer by increasing the convective heat transfer area. The optimization of heat sinks for single-phase immersion cooling is a topic of growing research interest [4-6]. Achieving an optimal design which utilizes natural convection inlet conditions can result in more efficient and cost-effective cooling solutions without the use of external pumping power. This decreases energy consumption and operating costs leading to enhanced system reliability and reduced energy consumption. Natural convection flow relies on buoyancy-induced flow of the dielectric liquid, driven by a density difference caused by the temperature difference between the heat sink and the fluid. This paper aims to optimize heat sinks for single-phase immersion cooling using natural convection.

In order to maximize the potential thermal performance benefits of single-phase immersion cooling, the heat sink must be optimized for the fluid properties of liquid cooling dielectric fluids. The general design optimization of parallel plate-fin heat sinks has been studied within the existing literature, primarily for heat transfer to air [7-8]. The primary research question of this study revolves around identifying the key design parameters that significantly impact the heat sink performance and determining the optimal configurations for minimizing heat sink thermal resistance specifically for immersion cooling natural convection.

This study proposes a CFD modeling methodology and multi-variable design optimization approach for determining the optimal fin characteristics for parallel plate-fin heat sinks for natural convection single-phase immersion cooling.

3.2 CFD Model Methodology

3.2.1 CFD Model Setup

A Computational Fluid Dynamics (CFD) model was developed in ANSYS Icepak consisting of a heat sink, thermal interface material (TIM), and 2D heat source as shown in Figure 3.1. A fluid domain cabinet was created to model the flow domain around the heat sink in natural convection single-phase immersion cooling.

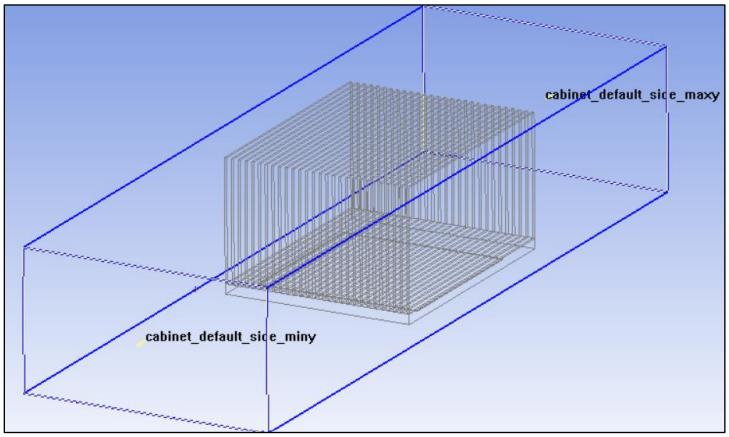

Figure 3.1 ANSYS IcePak Heat Sink Model

The model consists of both fixed and variable design parameters. For this study, the fluid domain is modeled as EC-110 dielectric fluid. The heat sink is modeled as an extruded aluminum parallel fin heat sink configuration with a fixed overall length of 110 mm, overall width of 86.1, base height of 4.5 mm, and overall height of 59.5 mm to model a 2U server heat sink configuration. The heater dimensions are modeled as 75 x 75 mm^2, based upon a review of current generation state-of-the-art server CPU sizes available.

Table 3.1 Fixed CFD Model Parameters

Model Parameter	**Value**
Heat Sink Base Thickness	4.5 mm
Heat Sink Overall Length	110 mm
Heat Sink Overall Width	86.1 mm
Heat Sink Overall Height	59.5 mm (2U)
Heat Sink Material	Al
CPU Heater Area Size	75 x 75 mm^2
Fluid Cabinet Length	350 mm
Fluid Cabinet Width	114.1 mm
Fluid Cabinet Height	63.5 mm
TIM Thickness	0.2 mm
Tim Thermal Conductivity	8 W/m-K
Fluid Inlet Condition	1.39e-05 kg/s
Fluid Outlet Condition	101,325 N/m^2

Computational Fluid Dynamics numerical modeling solves the Navier–Stokes equations of mass, momentum, species, and energy elementwise to calculate heat transfer in laminar flow conditions. For the current study, the flow is assumed to be laminar flow and therefore the transport equations for turbulence were not used. The following equations are solved elementwise to calculate laminar flow heat transfer:

Conservation of Mass:

$$\frac{\partial \rho}{\partial t} + \nabla \cdot (\rho \vec{v}) = 0 \tag{3.1}$$

For incompressible flow, (3.1) reduces to:

$$\nabla \cdot (\vec{v}) = 0 \tag{3.2}$$

Conservation of Momentum:

$$\rho \frac{\partial \vec{v}}{\partial t} + \rho(\vec{v} \cdot \nabla)\vec{v} = -\nabla p + \rho \vec{g} + \mu \nabla^2 \vec{v} \tag{3.3}$$

Conservation of Energy for the fluid domain:

$$\rho \frac{\partial h}{\partial t} + \rho \nabla \cdot (h\vec{v}) = -\nabla p + \nabla \cdot [(k + k_t)\nabla T] + S_h \tag{3.4}$$

Solid domain energy equation:

$$\frac{\partial}{\partial t}(\rho h) = \nabla \cdot (k \nabla T) + S_h \tag{3.5}$$

The Bousinessq approximative model is used to model a buoyancy-driven flow field by ignoring variable fluid densities except in the direction of a specified gravity vector. This approach treats fluid density as constant in the above equations, except for the momentum equation buoyancy term, which becomes:

$$(\rho - \rho_0)\vec{g} \approx -\rho_0 \beta (T - T_0)\vec{g} \tag{3.6}$$

Where $\rho 0$ is the constant fluid density, T_0 is the constant operating fluid temperature, and β is the fluid volumetric expansion coefficient.

3.2.2 CFD Model Validation

3.2.2.1 Mesh Sensitivity and Grid Independence

A mesh sensitivity analysis was performed considering a baseline heat sink design to verify analysis results are grid independent and mesh density and quality is sufficient. The default minimum element size generated by ANSYS Icepak is 1/20 of the specified lengthwise direction. For the conducted grid independence study, the model mesh was varied from coarse to fine by redefining a lengthwise mesh of 1/5, 1/20, 1/30, 1/40, 1/50, 1/60, and 1/100. Heat sink thermal resistance and heater temperature were calculated and tabulated for each mesh element count, shown in Figure 3.2.

Localized object parameter meshing control was utilized on the heatsink object within the Icepak model to ensure that a minimum of 5 elements are across the fin width, 15 elements along fin length, and 8 elements along fin height in all mesh density cases of the mesh sensitivity study. The junction temperature is within 3.2% of the value at 2,155,073 elements at 244,814 elements. Therefore, this overall mesh density was selected to conduct the optimization to significantly reduce computational workload.

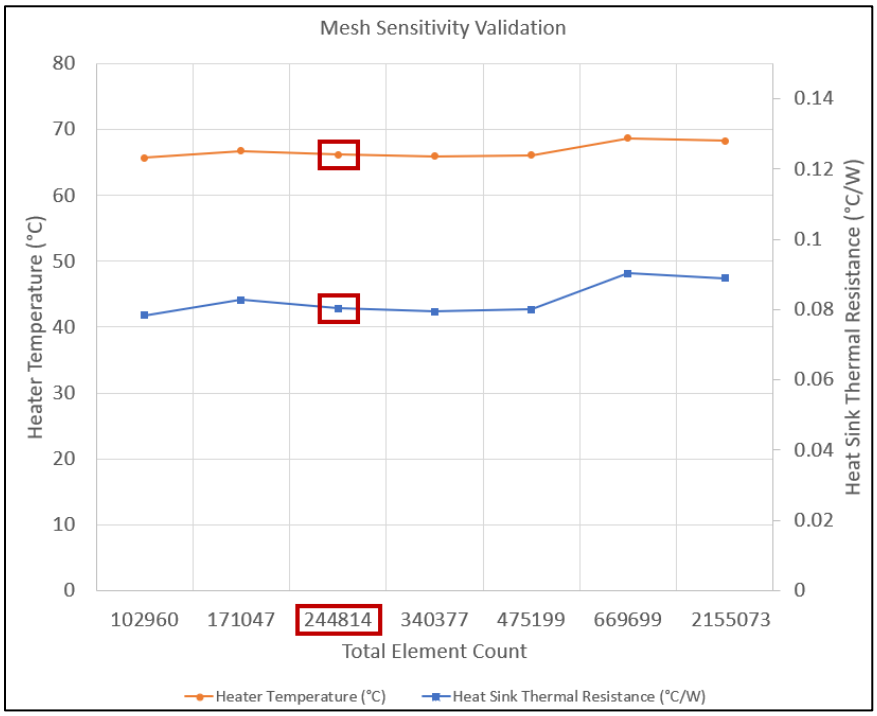

Figure 3.2 CFD Model Validation - Mesh Sensitivity Study

3.2.2.2 Natural Convection Modeling Approach

Validation of the fluid inlet flow conditions was carried out by determining the heater temperature for variable mass flow rate inlet conditions. This was conducted to determine the

validity of assuming that at low inlet flow rates, the flow may be assumed to be fully buoyancy-driven.

The mass flow inlet validation shows that the forced convection component is non-impactful to steady state heater temperatures below 0.01 LPM. Therefore, a mass flow rate fluid inlet condition of 0.001 LPM (1.39e-05 kg/s) was selected for the purpose of modeling buoyancy driven natural convection flow for the heat sink optimization process.

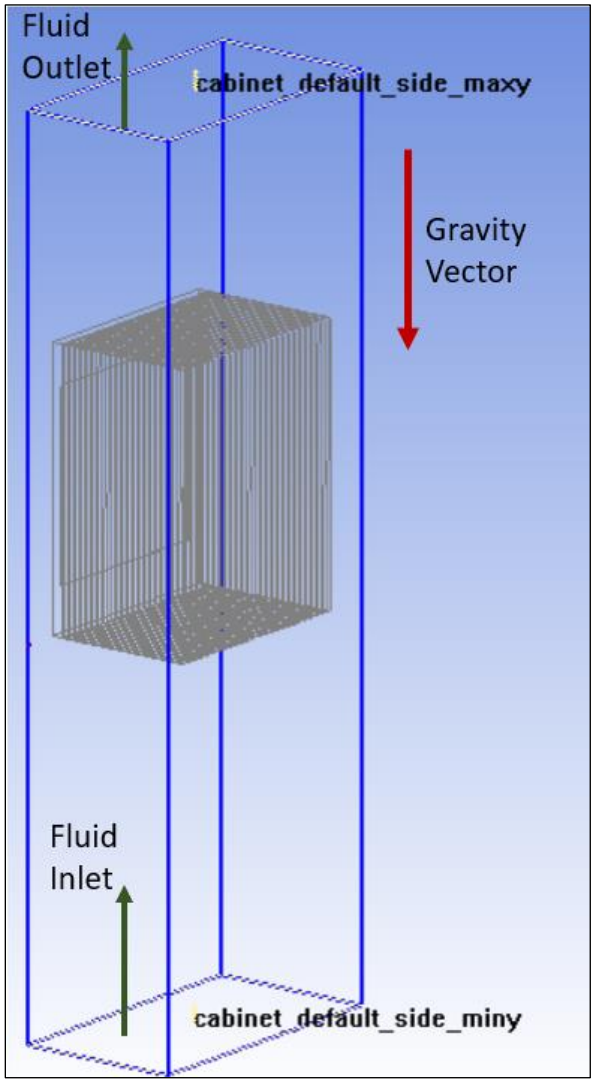

Figure 3.3 Natural Convection Boundary Condition

Table 3.2 Natural Convection Mass Flow Inlet Validation

Volumetric Flow Rate (LPM)	Mass Flow Rate (kg/s)	Heater Temp (°C)
0.00001	1.39E-07	72.1
0.0001	1.39E-06	72.1
0.001	**1.39E-05**	**72.1**
0.01	1.39E-04	72.1
0.1	1.39E-03	71.8
0.2	2.78E-03	71.4
0.4	5.55E-03	70.7
0.8	1.11E-02	69.1
1	1.39E-02	68.4

3.3 Heat Sink Optimization Methodology

ANSYS OptiSLang is a process integration and design optimization tool that enables automated parametric design studies within ANSYS solvers. OptiSLang allows for multi-objective and multi-design variables to be solved independently for a range of optimization parameter bounds. The parametric optimization tool utilized is an Adaptive Meta-Model of Optimal Prognosis (AMOP) sampling approach. AMOP sampling calculates a Coefficient of Prognosis (CoP) to determine predictive approximation quality of the model variables [9].

$$\text{CoP} = 1 - \frac{SS_E}{SS_T} \tag{3.7}$$

Where SST is the total model variation and SSE is the variation due to regression calculated as the sum of the square of prediction errors. Higher CoP values indicate more accurate data representation within the model, reduce output data postprocessing requirements, and allow for more direct assessment of design variable trends and surface response plots for the investigated design space. AMOP is more practical for optimization cases that contain a large number of input variables, and overcomes this hurdle faced by traditional meta-modeling approaches by assessing importance of each variable to the overall model and eliminating unimportant design variables.

The objective of this multi-variable optimization model is to determine the heat sink fin parameters that provide optimal thermal performance within natural convection single-phase immersion cooling. To achieve this, an AMOP optimization is carried out for the range of variable input parameters below for the described CFD model. The optimal heat sink fin design parameters, fin thickness and fin count, are determined for variable Heater TDP and Fluid Inlet Temperature.

Table 3.3 Variable Model Parameters

Model Parameter	Value
Heater TDP	250, 350, 450 W
Fluid Inlet Temperature	25, 35, 45 °C
Heat Sink Fin Thickness	0.6, 0.8, 1.0, 1.2, 1.4, 1.6 mm
Heat Sink Fin Count	20, 22, 24, 26, 28, 30
Objective Function	Minimize HS Thermal Resistance

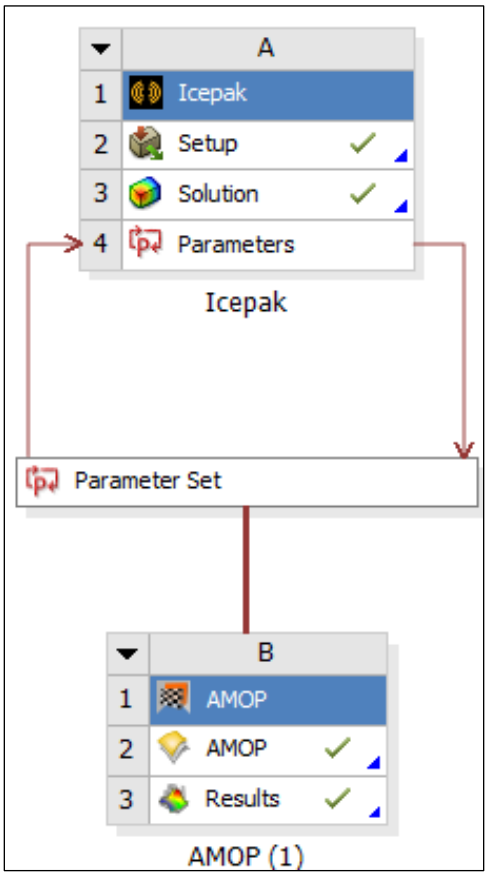

Figure 3.4 ANSYS Workbench optiSLang Integration

A design of experiments (DoE) is generated for the input design variables to conduct a sensitivity analysis for determining the effect of input variables on the objective function. This step is known as the design exploration phase where the design space is sampling using the selected sampling approach. In the current study, AMOP sampling is used to conduct the design space exploration. The outputs of heat sink thermal resistance and heater temperature are calculated for each sample and outputs from this phase are used to generate response surface plots and the total effects matrix in OptiSLang. The total effects matrix is used to determine the weighted relationship between input design variables and the output parameters and the objective function.

3.4 Results and Discussion

3.4.1 Optimization Variable Sensitivity Analysis Results

The first step in optimization data post-processing is assessing the sensitivity of the objective function and output variables to the input design variables through design exploration of the total effects matrix, to determine the relative impact of each design variable on the objective. Figure 3.5 displays the total effects matrix for the optimization sensitivity study. The objective function, heat sink thermal resistance, is primarily a function of Heater TDP, with a weight of 52.5%. In practice, this is often a hardware constraint rather than a tunable variable, with the goal of sufficiently cooling higher power dissipations. A key takeaway from the total effects matrix is that for the natural convection design space of this study, fin thickness has a significantly larger impact on heatsink thermal resistance than the number of fins, with a weight of 42.5% and 23.0%, respectively. Additionally, the total effects matrix is used to assess the total CoP of the model. The CoP is 100% for the objective function, and greater than 99% for all output variables. This

indicates that the data points generated by the AMOP sampling approach was successful and all input variables are meaningful to the outputs.

Figure 3.5 Multi-Design Variable Total Effects Matrix

3.4.2 Optimal Heat Sink Fin Parameter Results

Heat sink thermal resistance has been shown to be a function of both fin thickness (mm) and fin count from the total effects matrix (Figure 3.5). The dependent relationship between heat sink thermal resistance and fin design parameters can be visually represented with 3D response surfaces, shown in Figures 3.6 – 3.8 for various samples of Heater TDP and Inlet Fluid Temperature. The optimal heat sink fin thickness and fin count results for the design space are tabulated for each sampled TDP and Inlet Temperature in Table IV and V respectively.

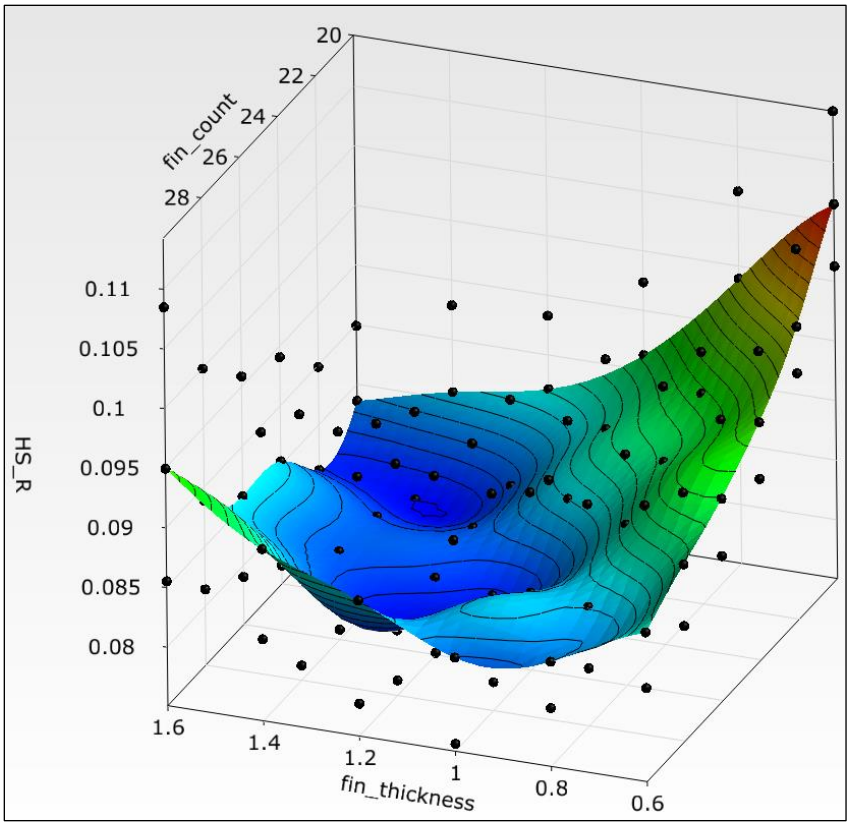

Figure 3.6 Optimization Response Surface Plot - Heat Sink Thermal Resistance for Fin Parameters at 25 C, 350 W

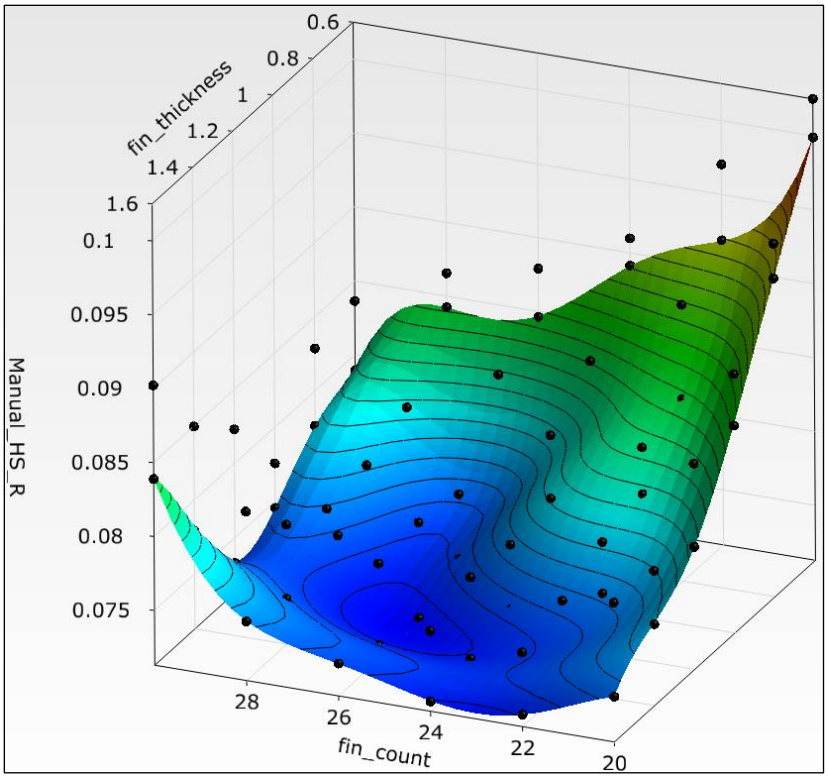

Figure 3.7 Optimization Response Surface Plot - Heat Sink Thermal Resistance for Fin Parameters at 35 C, 450 W

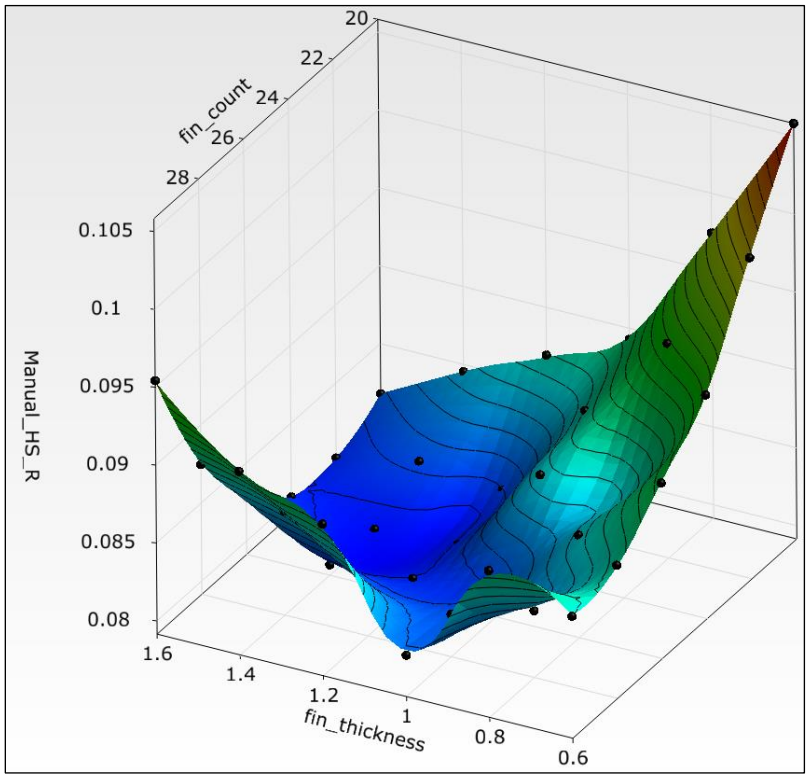

Figure 3.8 Optimization Response Surface Plot - Heat Sink Thermal Resistance for Fin Parameters at 45 C, 250 W

Table 3.4 Optimal Fin Thickness [mm]

Inlet Fluid Temperature	Heater TDP		
	250 W	350 W	450 W
25 °C	1.4	1.4	1.6
35 °C	1.4	1.6	1.2
45 °C	1.4	1.2	1.2

Table 3.5 Optimal Fin Count (Spacing [mm])

Inlet Fluid Temperature	Heater TDP		
	250 W	350 W	450 W
25 °C	22 (2.63)	22 (2.63)	22 (2.42)
35 °C	24 (2.28)	22 (2.42)	26 (2.20)
45 °C	24 (2.28)	26 (2.20)	26 (2.20)

3.5 Conclusions

Natural convection single-phase immersion cooling is of interest for its ability to cool higher power densities than forced-air cooling while maintaining a more overall energy-efficient thermal management system [10-11]. This study analyzed and optimized parallel fin heat sink design parameters for natural convection single-phase immersion cooling for various inlet fluid temperatures and chip power dissipations using ANSYS Icepak and OptiSLang for computational fluid dynamics simulations and multi-design variable sensitivity analysis and optimization. The

relative weighted impact of TDP, inlet fluid temperature, fin thickness, and fin count (spacing) were determined with an AMOP sampling of CFD data points and optimized for the variable problem constraints and boundary conditions.

Chapter 4

Machine Learning-Based Heat Sink Optimization Model for Single-Phase Immersion Cooling

Reprinted with Permission [30]

4.1 Introduction

The accelerating demand for data-driven technologies such as artificial intelligence, enterprise cybersecurity, and high-power edge computing is being met with rapid growth in the scale and number of data centers worldwide [1]. In order to meet critical performance requirements, data center computational workloads are rising while response times shorten. This necessitates electronic component device miniaturization and continual improvements in data center server density. Despite increases in IT equipment energy efficiency, the total energy consumption of data centers is projected to as much as triple over the next decade [2]. With power densities ranging 15-100 times greater than standard commercial buildings, effective thermal management of data centers is the key enabler for sustained future growth [3].

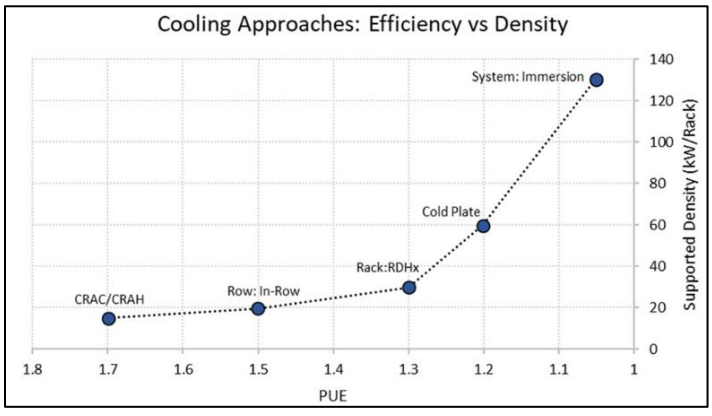

Figure 4.1 Comparison of Data Center Cooling Technologies [4]

Traditional air-cooling thermal management strategies are limited by the total thermal design power (TDP) that can be effectively cooled. To combat increased power-density and rising energy consumption, various propositions to either improve the operational performance of current cooling strategies [5-6] or the implementation of liquid-based cooling technologies [7-8] have been investigated. While there are multiple promising liquid-cooling methodologies, single-phase immersion cooling (SPIC) is established as a low infrastructure cost, low-complexity, and easy to deploy liquid-cooling approach. In addition to superior thermal performance over air-cooling, SPIC also provides reliability enhancements such as lower risk of system contamination and reduced vibration-induced component failure due to the complete submersion of the server in a dielectric fluid. The implementation of SPIC with data center servers initially designed for air-cooling applications requires additional design considerations, such as removal of air-cooling hardware, hermetically sealing open drives, and material compatibility must be addressed.

Air-cooled component heat sinks are optimized for the fluid properties of air as the cooling fluid. Thus, in order to take full advantage of the improved thermal properties of SPIC with air-cooled hardware, the component heatsinks on high power-density electronics such as the central processing unit (CPU) and the graphics processing unit (GPU) must be optimized for the SPIC dielectric fluid. The thermal performance optimization of plate-fin heat sinks has been broadly studied in the literature [9-12].

4.2 Overview of Machine Learning Modeling Approach and Techniques

4.2.1 Supervised Machine Learning Overview

The emergence of machine learning (ML) algorithms for data-driven statistical modeling as an effective new predictive tool for modeling the behavior of complex non-linear systems allows for new advances not easily achieved through traditional physics-based modeling approaches. Given a large enough data set which appropriately defines the problem space, ML techniques have been used to predict system behavior in a broad range of fields in engineering, medicine, and the financial sector [13-16]. More specifically, ML has seen a growing utilization in Data Center thermal performance evaluation [17-19], energy efficiency enhancement [20-21], and thermal management infrastructure control[22-24]. The current state of the literature focuses on macro scale thermal management strategies, and to the best of the author's knowledge, no current literature studies the applications of ML on server or heat sink level performance optimization. The ML algorithms analyzed in this study such as polynomial regression, random forest, and neural networks are also currently being utilized to build complex system models within other thermal science and thermal management applications [25-26].

Supervised machine learning algorithms require a labeled training data set $\{\boldsymbol{x_i},\boldsymbol{y_i}\}$,$i$=1,2,…,$n$, $\boldsymbol{x_i} \in \mathbb{R}^d$, and $\boldsymbol{y_i} \in \mathbb{R}^k$.

The learning task is to define an unknown function f: $\mathbb{R}^d \rightarrow \mathbb{R}^k$, such that:

$y_i = f(x_i) + \delta_i$, $\forall\, i \in \{1,2,\ldots,n\}$ where δ_i is the error associated with the i^{th} data point. In the context of this study, $\boldsymbol{x_i}$ is the input feature vector containing five input features: i. heat sink height, ii. fin thickness, iii. number of fins, iv. fluid flow condition, v. heat sink material.

The following ML techniques differ in their approach to obtain the target function and minimizing regression error.

Two primary considerations when selecting a supervised machine learning algorithm for implementation are accuracy of the regression model and model interpretability. The model's prediction accuracy is characterized by the model's flexibility to fit the dataset. For statistical regression-based models, accuracy is inversely proportional to the amount of underlying non-linearity in the data set due to the inherent rigidness of the technique. However, statistical models exhibit high interpretability by outputting a closed-form mathematical equation describing the dependency between input features and target outputs. Non-linear ML methods, such as random forest regression and neural networks, are very effective at modeling non-linear datasets, but sacrifice model interpretability and are therefore described as very accurate black-box models.

4.2.2 Polynomial Regression

Statistical regression fits a multi-dimensional hyperplane onto the data point domain such that the summation of the squared distances from the hyperplane to each data point is minimized. The flexibility of the hyperplane is represented by the order of equations used to build the model. It was determined through an exhaustive search of polynomial order that for the present data set, a polynomial regression of 3rdorder represents the best-fit of the data and represents the relationship between the input features and target outputs. A 3rdorder polynomial equation was determined to be the best predictor of outputs without risking overfitting the model to the training data, and adding higher order terms did not produce any additional benefit in the model's predictive performance capability.

4.2.3 Random Forest

Decision trees are based upon a set of splitting rules which are used to stratify the predictor space into simplified regions in order to optimize the objective function. Because the splitting rule set which divides the predictor space is graphically drawn as an upside down tree with mean observations in the leaves, or nodes, at the bottom and root at the top, this stratification is denoted as decision tree-based methodology. In particular, the random forest ML algorithm takes advantage of the bagging technique, as well as random feature selection, providing a more established ensemble decision tree. A single decision tree model has been shown to have high variance, since it overfits the data, performing with very high accuracy on the training data, but fails to generalize to unseen test data. One method for overcoming this inherent overfitting seen in decision trees is by utilizing the bagging technique (or bootstrap aggregating). Bagging reduces variance and produces a model that may successfully make generalized predictions on unseen test data. Bagging fits n number of decision tree models to n data samples and the average of all trees is taken as the final model used for predicting outputs. Finally, by considering only a subset of the features at each decision split for each internal node, random forest algorithms introduce an additional improvement over bagged decision trees in that they de-correlate each individual tree within the forest.

4.2.4 Neural Network

Neural networks (or Artificial Neural Networks) are distributed adaptive machines comprised of computing cells, or neurons, that are capable of storing acquired knowledge through experiential training and learning processes. The intercell, or interneuron connections are associated with synaptic weights which are used to store experiential knowledge. Typically, neural networks are organized with neuron layers. The first layer, or input layer, is responsible for receiving the training data input features. Successive layers, or hidden layers, are passed the outputs from the preceding layer. The number of layers within the neural network and number of synaptic weights contained within each hidden layer are tunable features of the neural network algorithm and have a great amount of impact on the final accuracy of the model produced. The final hidden layer, or output layer, outputs the predicted target values.

4.3 Machine Learning Model Implementation

4.3.1 Machine Learning Dataset

The primary objective of this study is to investigate the efficacy of machine learning models to predict thermal resistance and pressure drop for a parallel plate-fin heat sink of an air-cooled Open Compute server in order to optimize performance for immersion cooling applications. This study utilizes results obtained from a previous work in which an experimentally validated high-fidelity computational fluid dynamics (CFD) ANSYS model of the server was used to conduct a multi-objective and multi-design variable heat sink optimization study in

OptiSLang[27].OptiSLang is a design optimization extension utilized as part of ANSYS Workbench, and can be directly integrated with any ANSYS thermal, structural, electrical, or fluid tools. The optimization design variable input parameters are: i. heat sink height, ii. fin thickness, iii. number of fins.

Fins allow for increased overall heat sink surface area, improving the overall heat transfer of the heat sink by lowering thermal resistance. However, a competing performance metric is pressure drop, whereby the heatsink impedes the continuity of the cooling fluid. Optimal fin geometry minimizes both thermal resistance and pressure drop and is a function of multiple factors, such as fluid inlet flow rate and density. Each design variable contained a range of 6 discrete values, and included two flow boundary conditions, forced flow inlet at 2 LPM, and natural convection.

Each heat sink configuration and flow condition were carried out for aluminum and copper heat sink materials, and peak utilization power for each CPU is 115 W. The baseline air-cooled heat sink parameters are as follows: fin count of 41, base thickness of 4 mm, fin thickness of 0.23 mm, with a fin height of 37mm. The dielectric fluid selected for the CFD study is a commercially available synthetic fluid, EC100, and is held as a constant throughout the analysis.

Through these four parametric simulations a database of 864 data points for outputs of thermal resistance and pressure drop were generated and used to train and test three machine learning models; polynomial regression, random forest regression, and neural network regression in their ability to predict these two output parameters as a function of the five inputs.

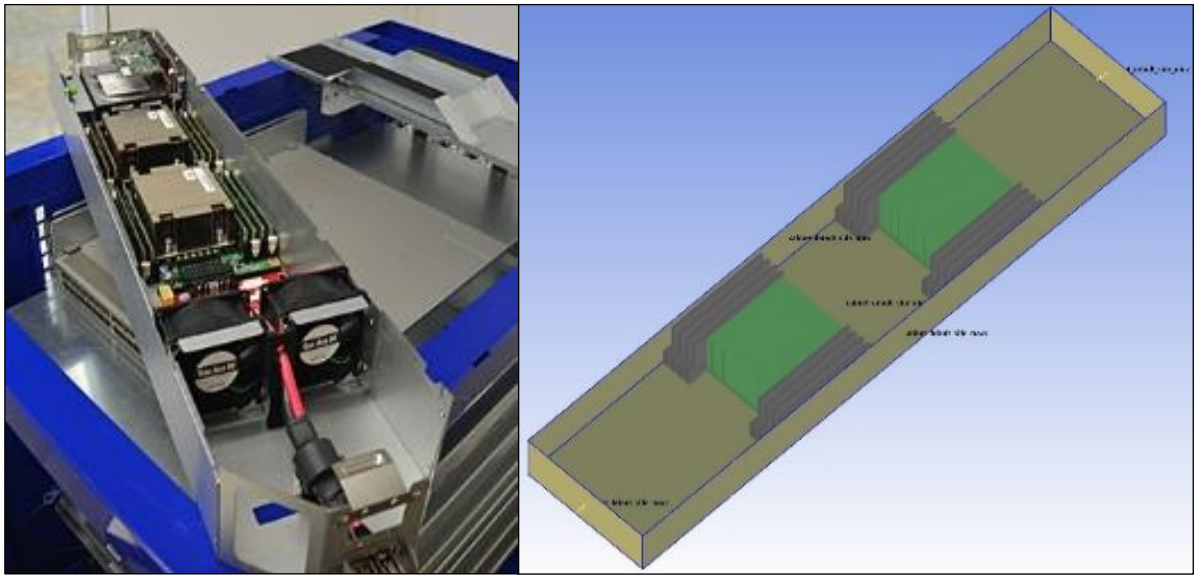

Figure 4.2 Open Compute Air-Cooled Server and ANSYS Model

Table 4.1 Optimization Input Design Variables (Bold are Air-Cooled Baseline Values)

	Overall Heat Sink Height (mm)	Fin Thickness (mm)	No. of Fins	Material	Fluid Flow
	26	***0.23***	25	Aluminum	0 LPM
	29	0.32	27	Copper	2 LPM
	32	0.41	29		
	35	0.5	31		
	38	0.59	33		
	41	0.68	***35***		
Step Size	3	0.09	2		
Design Points	6	6	6	2	2

Total Number of Design Points (6 x 6 x 6 x 2 x 2) = 864

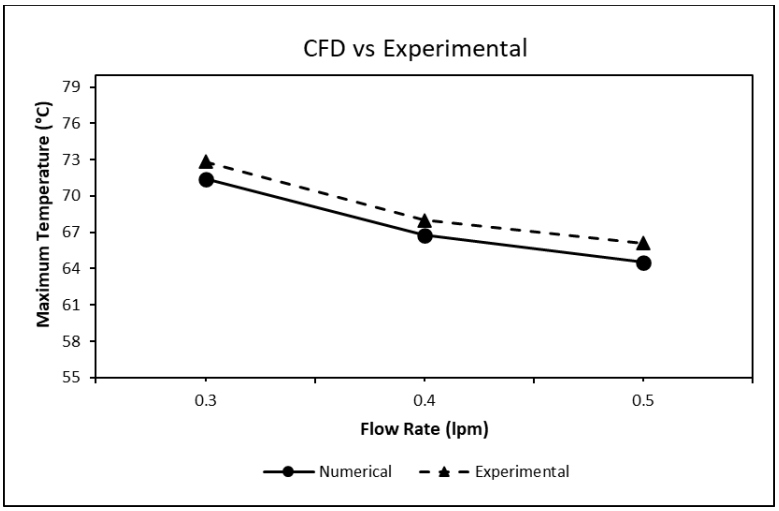

Figure 4.3 Comparison of Junction Temperature CFD Simulation Results and Experimental Validation Data [27]

This study analyzes three supervised learning algorithms and compares their ability to build accurate predictive models to map heat sink design features and boundary conditions as inputs to two discrete outputs; heat sink thermal resistance and heat sink pressure drop. The objective of this comparison is to identify interesting features of the algorithms, analyze hyperparameter tuning sensitivity, and compare the evaluated predictive capabilities and performance between algorithms.

The training dataset was generated by the numerical CFD simulations as discussed in Section 2. 216 data points of heat sink fin design features were simulated for both aluminum and copper heat sink material, and for forced and natural convection single-phase immersion cooling inlet flow boundary conditions. These 4 parametric simulations of 216 design points generated a total of 864 data points, which were used to develop three predictive models build by training supervised learning algorithms, polynomial regression (PR), random forest (RF), and neural

network (NN). During data preprocessing, 7 data points were removed due to result tabulation error, resulting in a total of 857 data points. The dataset was divided by allocating 90% of the total data points to the training data set, with the remaining 10% allocated to the test data set to evaluate the trained model's predictive accuracy. The training set of 771 data points was further subdivided with 10% of that allocated to a cross-validation data subset. The input feature vectors were standardized prior to model training. The target output vector for heat sink thermal resistance was passed through an elementwise multiplier of 100 in order to obtain the same order of magnitude for both thermal resistance and pressure drop. This was done in order to eliminate biasing the model weights toward the target with larger scaler values. Once the model weights were determined, the target outputs for thermal resistance were passed through an elementwise divider of 100 to return the final predicted values for thermal resistance.

4.3.2 Hyperparameter Tuning and Cross Validation

Overall model performance and prediction accuracy require that algorithm hyperparameters are tuned correctly in order to define model complexity and sufficient capacity for non-linearity of the dataset. In this study, hyperparameters considered for decision tree-based random forest regression include number of trees in the forest, maximum tree depth, split sampling quantity, and minimum data points allowed in a leaf node. Hyperparameters considered for neural network were number of hidden layers and dimension (number of nodes per layer) as the structural parameters. It should be noted that the chaining of linear (dense) layers is arbitrary and provides

no additional representative power, therefore a rectified linear unit (ReLU) activation function was utilized following each hidden layer to model non-linearities, and the Adam Optimizer was used for weight updates.

Hyperparameter optimization was carried out through exhaustive grid-search cross-validation (CV) with stratified 10-folds as it is believed that this would show better representation of the samples within the dataset. During training, learning curves and validation curves were plotted using averages that represent averages for accuracy across the folds; standard deviation bars represent the standard deviation bound of across the folds (+/-).

4.3.3 Metrics of Performance

Algorithm Metrics of Performance (MOP) were defined up front to establish a common comparison between the effectiveness of each learning algorithm when fitting the data. Mean Absolute Error (MAE) was used for learning curves and comparison metrics since it maintains the same units as the outputs being predicted and therefor can be directly compared. Mean Absolute Error (MAE) is calculated by the summation of the absolute difference between the predicted output value and actual output value, divided by the total number of data points in the set.

$$\mathrm{MAE} = \frac{\sum_{i=1}^{n} |\boldsymbol{y_i} - \hat{y}_i|}{n} \tag{4.1}$$

4.3.4 Supervised Learning Analysis Framework

A Python framework was designed to train algorithms, taking advantage of existing machine learning libraries: scikit-learn for polynomial regression, random forest regression and validation curve/learning curve/grid search functionality while tensorflow was leveraged for neural network development. A keras API was used to combine tensorflow functionality with the same scikit-learn tools used to analyze the scikit-learn based learning algorithms [28-29].

4.4 Results and Discussion

4.4.1 Machine Learning Model Results

4.4.1.1 Learning Curve Comparison

Learning curves provide insight into how effective a particular model is when exposed to increasing quantities of training data. The scikit-learn function “learning_curve” was used to generate 10-fold cross-validation models and statistics were plotted to examine model bias and variance. At each training data size, the mean of the training and validation MAE over all 10-fold model fits was computed. The standard division of the MAE for training and validation was also computed over all 10-folds to provide information about each model’s level of consistency when fit to different subsets of the training data (as the amount of training data increases). The +/-1 standard deviation bound is shaded on each learning curve.

An initial sweep over polynomial degrees showed lowest validation error for a degree 3 polynomial regression model. The random forest regressor and neural network were tuned

similarly to yield the results shown in Figures 5 and 6. From Figure 4, the polynomial regression algorithm MAE does not change significantly with each fold and when more data is added the model; the model improves with more data but begins to level off at 693 samples. The polynomial regression algorithm also shows the lowest variance gap of the three algorithms, the difference between the training and cross-validation curves which reflects consistency in the model's ability to generalize to a similar level of error as the model's error on the training data.

Between the low variance gap and convergence of the learning curves, the polynomial regression algorithm would not be expected to improve significantly if additional data was added. The random forest regressor learning curve (Figure 5) shows less certainty in the model given different folds as the cross-validation standard deviation is wider; this shows that the MAE varied more across each of the 10 folds when compared to polynomial regression. Despite being a more complex algorithm and showing signs of potential improvement with additional data (positive slope at 693 samples), the random forest regressor does not achieve the same desirable learning curve outcomes as that of the polynomial regression algorithm: low variance gap between training and cross-validation error. The neural network performed better than the random forest regressor in terms of training and cross-validation error curves and showed convergence of close errors over both. The neural network architecture chosen (each hidden layer densely connected with a ReLU after each) converged to a lower score than that of polynomial regression.

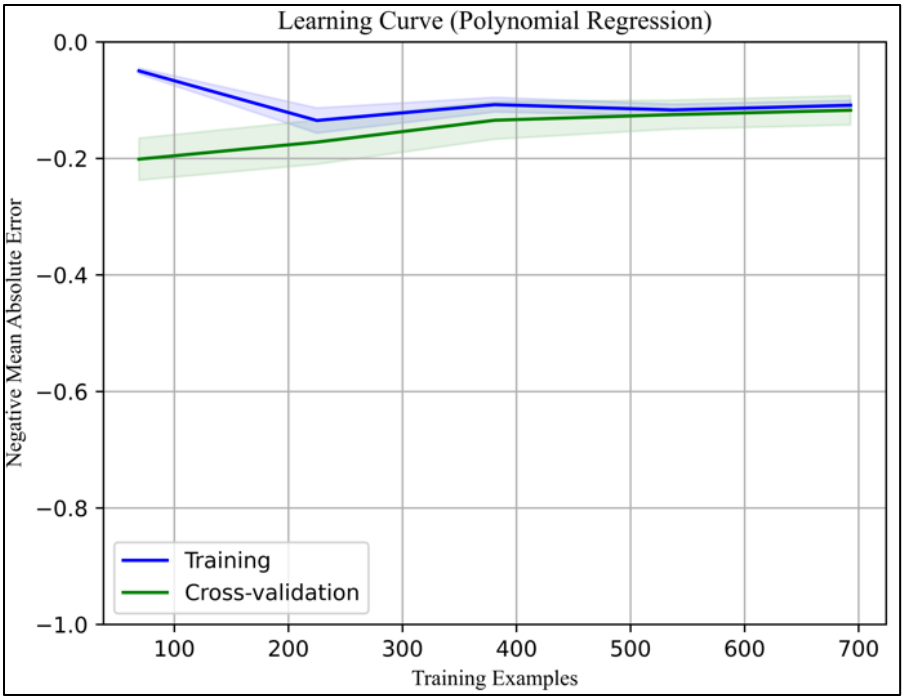

Figure 4.4 Polynomial Regression Model Learning Curve (Polynomial Degree = 3)

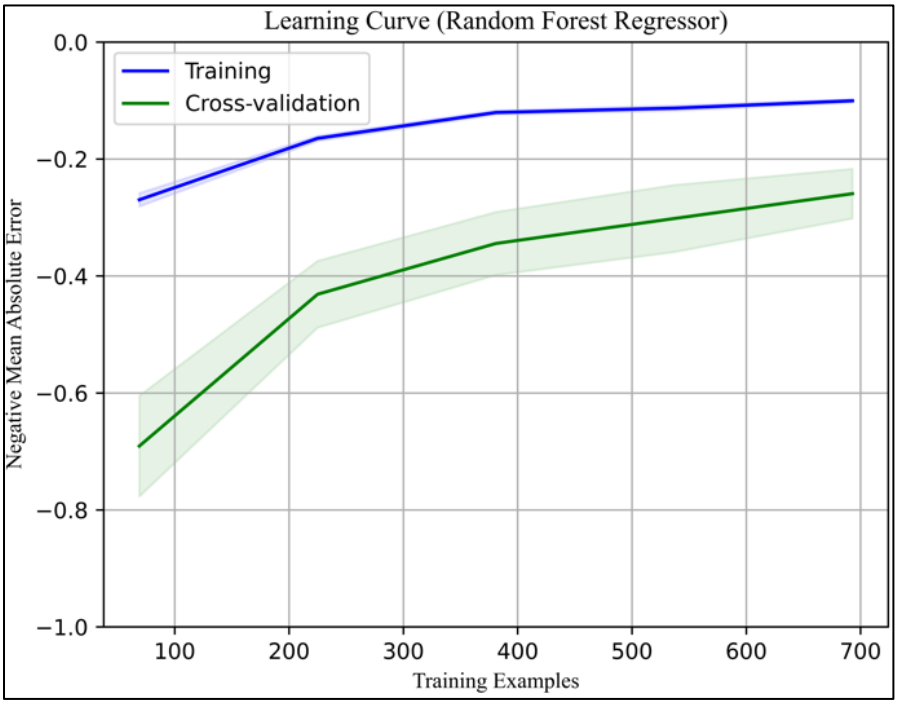

Figure 4.5 Random Forest Regressor Model Learning Curve (Estimator Count = 150, Tree Maximum Depth = 10)

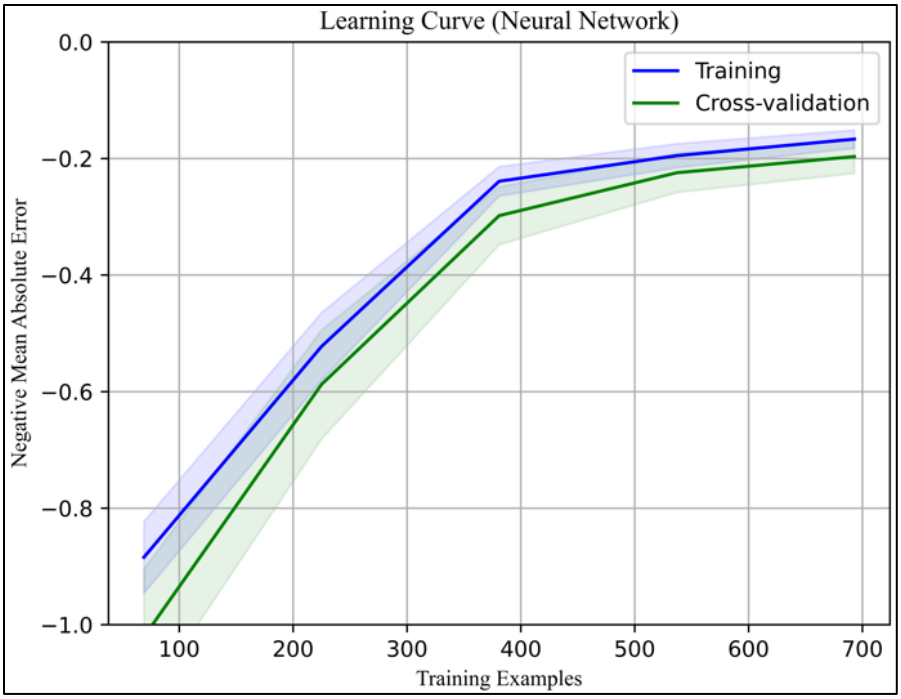

Figure 4.6 Neural Network Model Learning Curve (Hidden Layers = 4, Hidden Layer Dimension = 128, Learning Rate = 0.0005, Epochs = 50, Batch Size = 32)

4.4.1.2 Thermal Resistance Prediction Comparison

Table II. tabulates the mean absolute error (MAE) in predicted Thermal Resistance (°C/W) and Pressure Drop (Pa) compared to CFD computationally computed values for each algorithm's training set, test set, and total cumulative performance as quantified by MAE.

For thermal resistance, polynomial regression performed best, with a MAE of 0.00052 °C/W (training data), 0.00063 °C/W (test data), and total overall MAE of 0.00053 °C/W. Intuitively, this means that on average, given a set of inputs within the defined problem space, the model will predict thermal resistance +/- 0.00053 °C/W. Random forest performed with a total overall MAE of 0.000875 °C/W, and Neural Network performed with a total overall MAE of 0.00123 °C/W.

Notably, Polynomial Regression performing better than RF and NN indicates that while the governing physics of fluid mechanics and heat transfer are non-linear, the author's explanation for this result is that inlet fluid properties that are primarily responsible for the non-linear behavior, including temperature and flow rate, are held constant for the data set used to train the models. Given this, it is expected that introducing more data for variable fluid inlet temperature and flow rate would result in a more non-linear data set which would benefit further from the complex non-linear machine learning algorithms.

Unity plots (Figures 7-9) were used to verify the outputs of the model relative to the CFD values for both the training and test data. The test data was withheld until final model selection and never used as part of the training or incremental validation processes. All three models fit consistently across thermal resistance values with the exception for a distinct few training and test outliers seen in the same data range for each model.

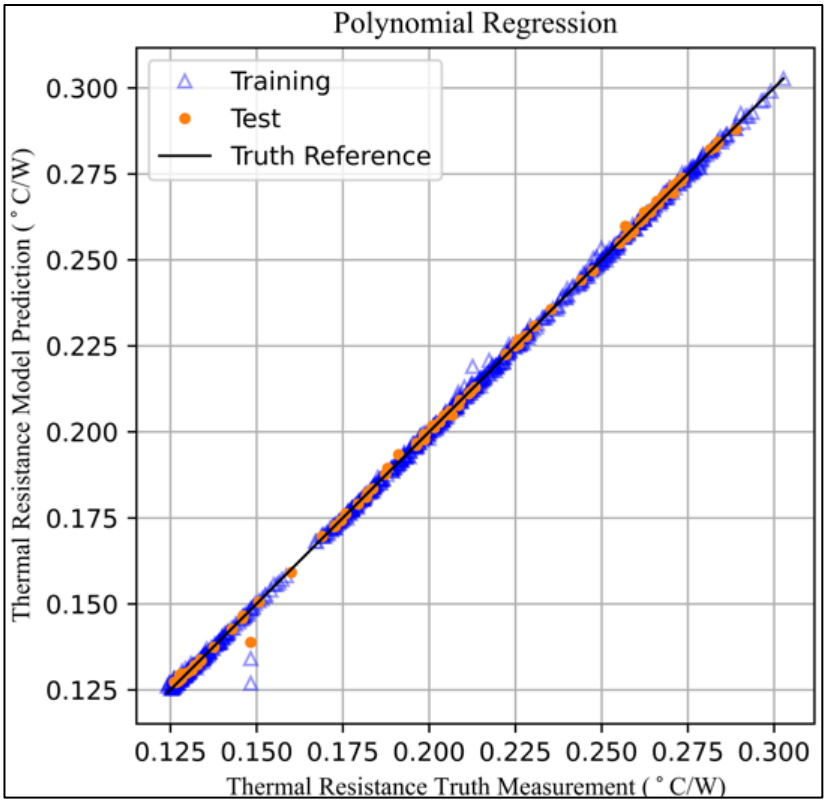

Figure 4.7 Polynomial Regression Model Unity Plot for Thermal Resistance (Polynomial Degree = 3)

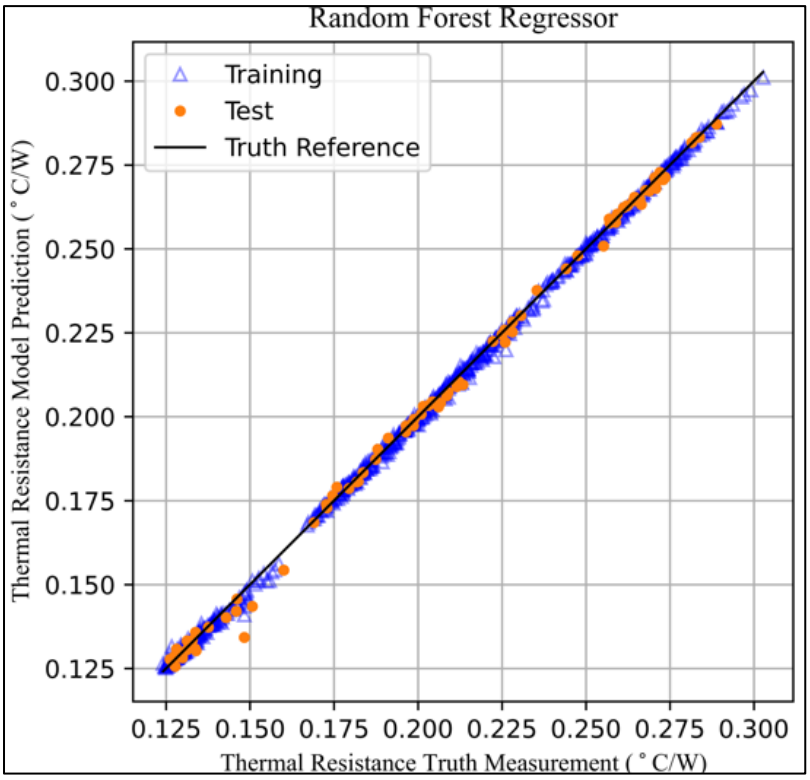

Figure 4.8 Random Forest Regressor Model Unity Plot for Thermal Resistance (Estimator Count = 150, Tree Maximum Depth = 10)

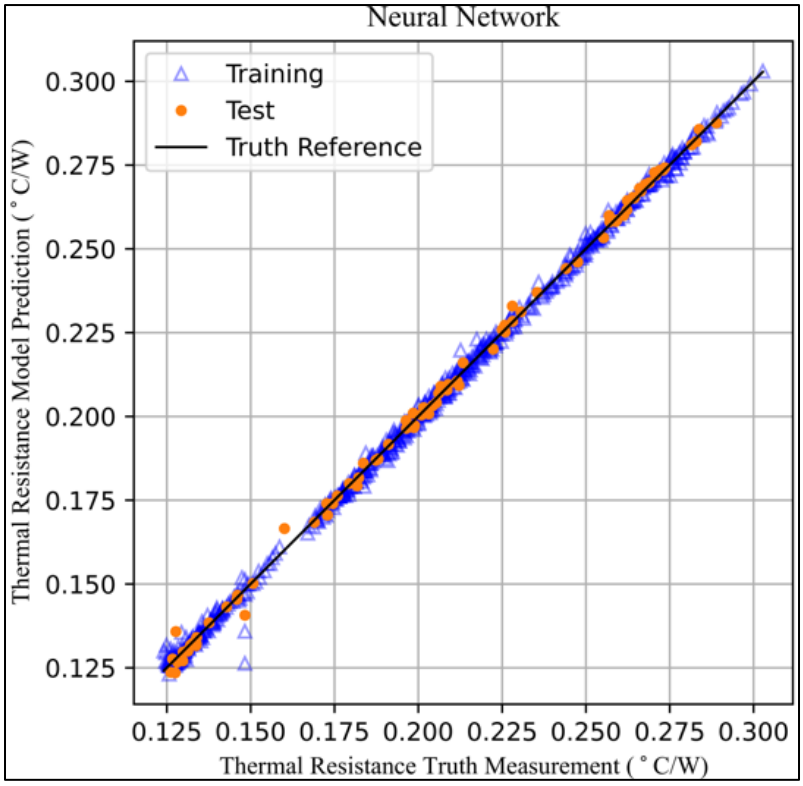

Figure 4.9 Neural Network Model Unity Plot for Thermal Resistance (Hidden Layers = 4, Hidden Layer Dimension = 128, Learning Rate = 0.0005, Epochs = 50, Batch Size = 32)

4.4.1.3 Pressure Drop Prediction Comparison

Given that heat sink pressure drop is strongly dependent upon fluid velocity, the resulting pressure drop is strongly grouped based on flow regime. Natural convection flow conditions resulted in lower pressure drops overall compared to forced flow. Figures 10-12 show pressure drop results for all data points, while Figures 13-15 display only those results associated with the natural convection inlet flow condition.

Unity plots for pressure drop showed that each model underestimated consistently at higher pressure drop values. The outliers are also more extreme (Figure 10), which shows there are particular inputs that all three algorithms have a difficult time predicting. The clustering of the

lower pressure drop values required isolation of the low region (0 to 2.25 Pa) on a separate plot for each algorithm. Interestingly, despite having higher training and test error, the random forest regressor fit the lower pressure drop values much tighter than the neural network and polynomial regression algorithms. This is likely because the random forest averages predictions over 150 estimators and is therefore less susceptible to variance.

Table 4.2 Performance Summary of Machine Learning-Based Predictive Models on Thermal Resistance and Pressure Drop

Mean Abs. Error	Polynomial Regression			Random Forest			Neural Network		
	Train	Test	Total	Train	Test	Total	Train	Test	Total
Thermal Resistance [°C/W]	0.00052	0.00063	0.000531	0.00079	0.00164	0.000875	0.0012	0.00149	0.00123
Pressure Drop [Pa]	0.168	0.274	0.1786	0.131	0.411	0.159	0.200	0.356	0.216

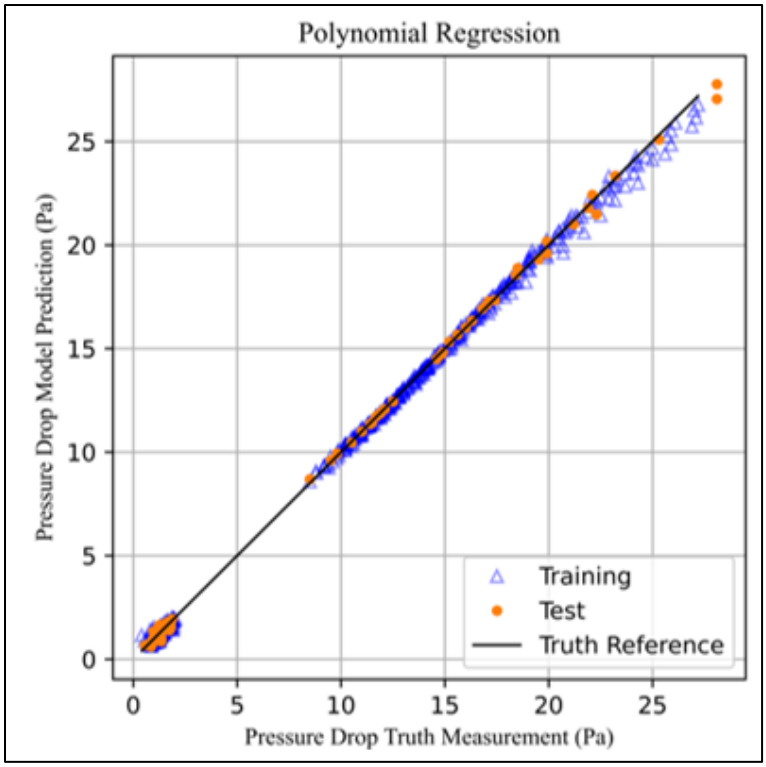

Figure 4.10 Polynomial Regression Model Unity Plot for Pressure Drop (Polynomial Degree = 3)

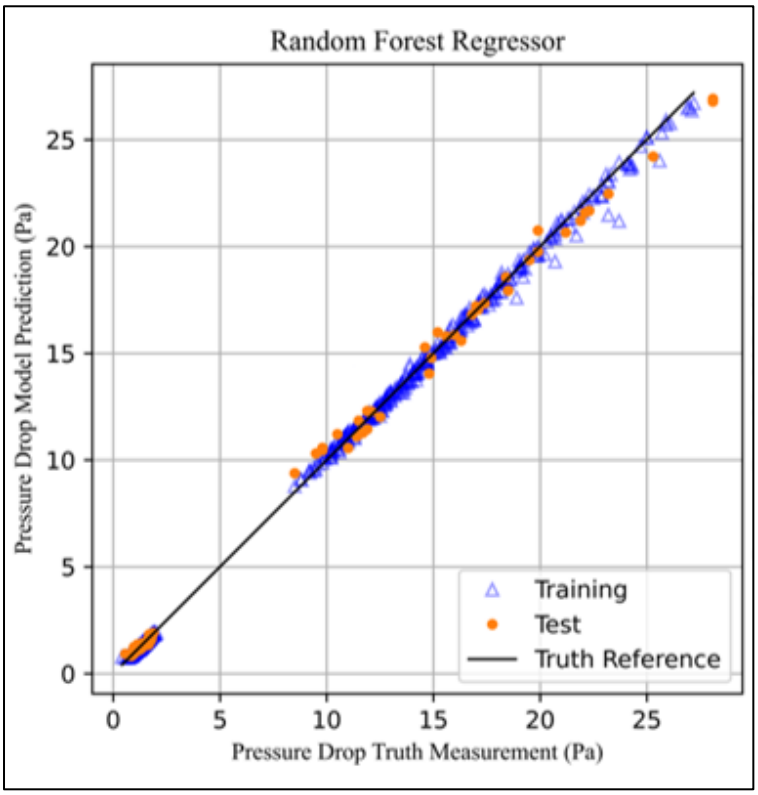

Figure 4.11 Random Forest Regressor Model Unity Plot for Pressure Drop (Estimator Count = 150, Tree Maximum Depth = 10)

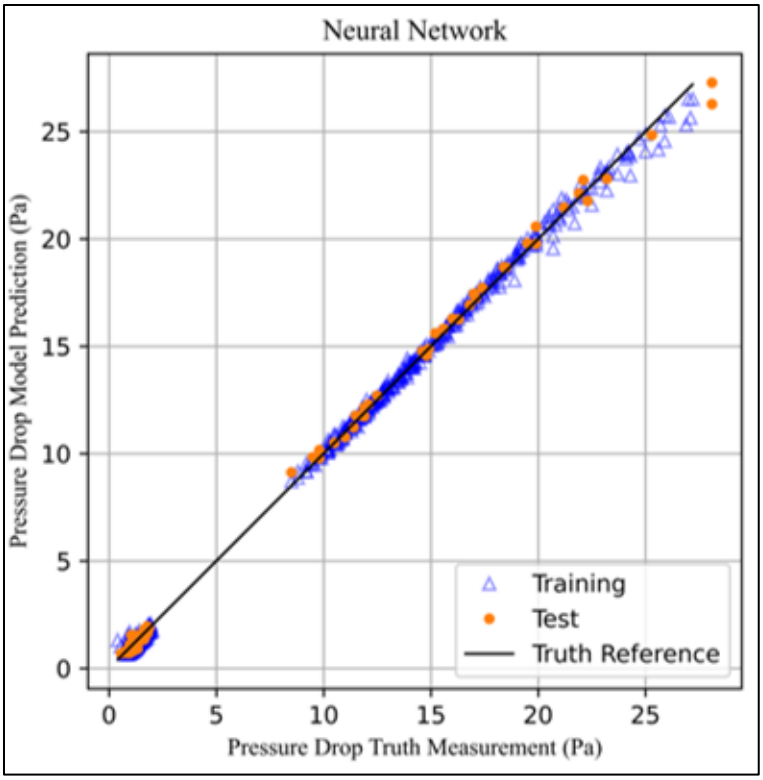

Figure 4.12 Neural Network Model Unity Plot for Pressure Drop (Hidden Layers = 4, Hidden Layer Dimension = 128, Learning Rate = 0.0005, Epochs = 50, Batch Size = 32)

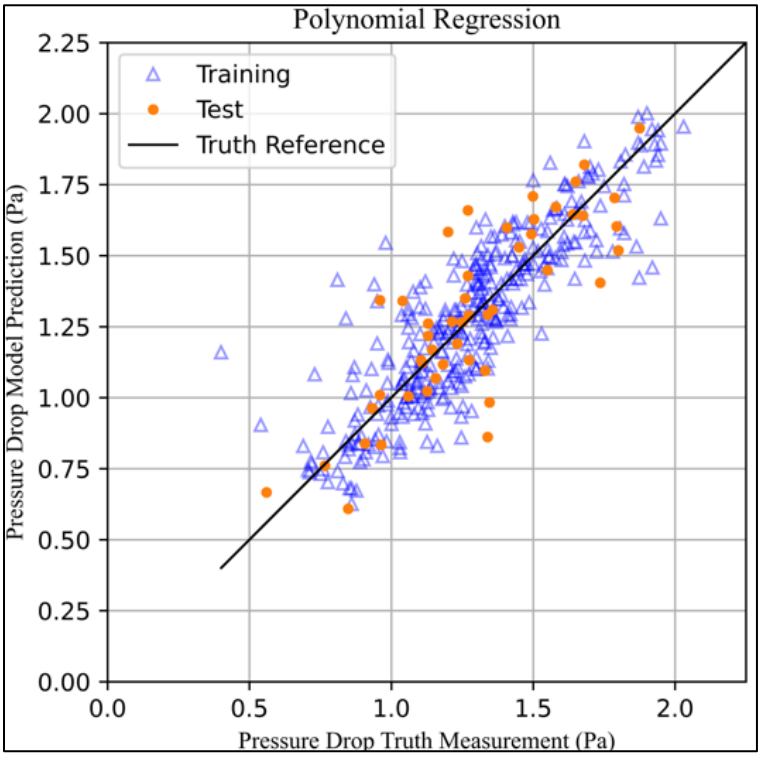

Figure 4.13 Polynomial Regression Model Unity Plot for Natural Convection (Polynomial Degree = 3)

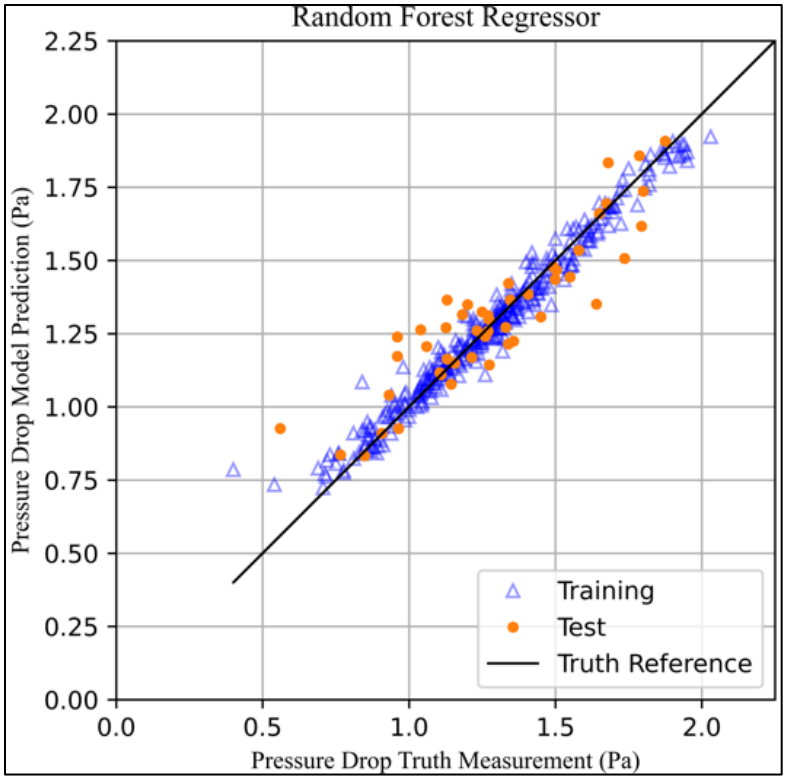

Figure 4.14 Random Forest Regressor Model Unity Plot for Natural Convection (Estimator Count = 150, Tree Maximum Depth = 10)

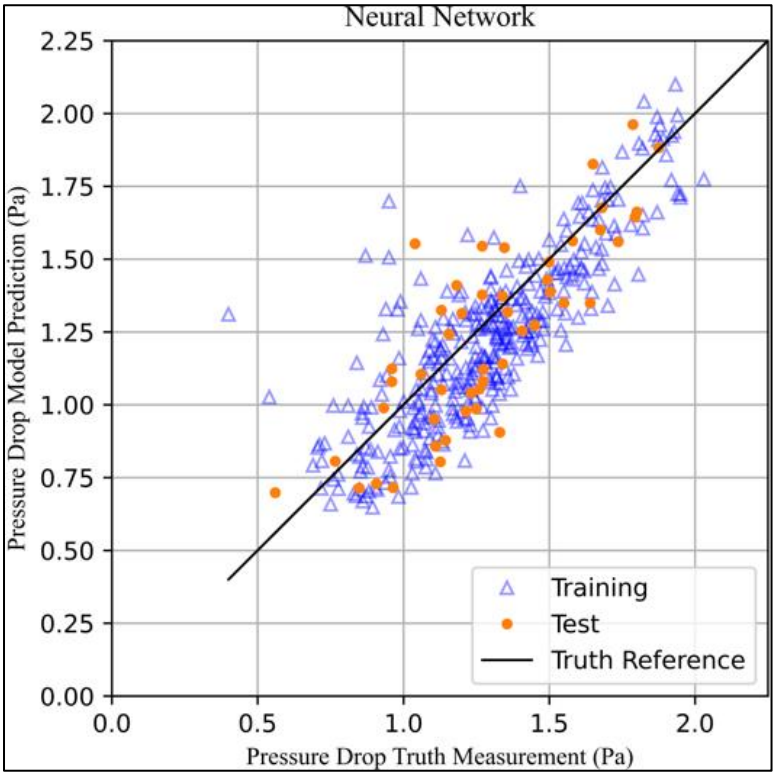

Figure 4.15 Neural Network Model Unity Plot for Natural Convection (Hidden Layers = 4, Hidden Layer Dimension = 128, Learning Rate = 0.0005, Epochs = 50, Batch Size = 32)

4.5 Conclusion

This study demonstrates the predictive utility of machine learning algorithms to model multiple thermal performance parameters for varying data center server heat sink design inputs. Using steady state multi-objective and multi-design variable optimization simulation data, a databank of 864 data points is used to train and test ML predictive models for heat sink thermal resistance and pressure drop. The methodology of developing and utilizing numerical simulation data in support of training machine learning models as shown in this work can be applied to various thermal management analyses to identify and optimize parameters to maximize thermal performance. While this study successfully highlights the premise of utilizing machine learning approaches for modeling server level data center thermal management, expanding the design space for which the model is applicable by including additional variables such as inlet flow velocities, inlet fluid temperatures, or chip power dissipation would further broaden the applicability of the model.

Chapter 5

Future Work

Future work may be done to further refine the parameters of the Machine Learning (ML) algorithms to enhance predictive accuracy. This refinement process will involve fine-tuning the algorithmic inputs and optimizing their interaction to yield more reliable and precise predictions for multiple output variables within the single—phase immersion cooling design space.

Another important aspect of future work involves expanding the model design space to incorporate additional design variables, thus increasing the model's predictive capability. The design space should be broadened to include various input variables such as heatsink design parameters, dielectric fluid properties, inlet fluid temperature, inlet fluid flow rate, chip power dissipation, and chip base area. By integrating these inputs, the model will be better equipped to predict key output variables, including heatsink thermal resistance, heatsink pressure drop, and junction temperature. This expanded model will provide a more comprehensive understanding of the thermal behavior in advanced cooling systems.

Additionally, future applications under investigation are focused on developing a machine learning-based design tool for single-phase immersion cooling systems. This tool aims to assist in the design and optimization of such systems by leveraging ML models to predict and evaluate performance metrics. Moreover, another critical application under consideration is the development of real-time controllers for single-phase immersion cooling. These controllers wibe based on live system modeling and will enable dynamic adjustment of cooling parameters to maintain optimal thermal performance under varying operational conditions. These future developments will contribute significantly to the advancement of thermal management solutions in high-performance computing systems.

References

5.1 Chapter 1 References

[1] Electric Power Research Institute. (2024). “Powering Intelligence; Analyzing Artificial Intelligence and Data Center Energy Consumption”. Research Whitepaper Report Number 000000003002028905

[2] Eric Masanet and Nuoa Lei. (2020). “How much energy do data centers really use?”. McCormick School of Engineering and Applied Science, Northwestern University

[3] TechSparks (2023). “Is Smaller Always Better for Transistor Size?” https://www.tech-sparks.com/size-of-transistors/

[4] Data Center Frontier (2022). “Sturdier Servers: Cloud Platforms Say Servers Living Longer, Saving Billions” https://www.datacenterfrontier.com/cloud/article/11427600/sturdier-servers-cloud-platforms-say-servers-living-longer-saving-billions

[5] MiTAC Computing Technologies (2024). “OCPserver E7278-S (E7278) (EOL)” https://www.mitacmct.com/OCPserver_E7278_E7278-S

[6] Ampere Computing (2024). “2U Mt. Jade 2S” https://amperecomputing.com/systems/altra/2u-mt-jade-2s-nvme

[7] Sivaraju, K. B., Saini, S., Simon, V., Bansode, P., Lamotte-Dawaghreh, J., Gupta, G., Herring, J., et al. (2022). “Comparative Study of Single-Phase Immersion Cooled Two Socket Server in Tank and Sled Configurations.” ASME 2022 International Technical Conference and Exhibition on Packaging and Integration of Electronic and Photonic Microsystems

[8] Y. Wang, et al., "Investigation on Immersion Cooling Solution for Hyper-scale Data Center Application," 2023 22nd IEEE Intersociety Conference on Thermal and Thermomechanical Phenomena in Electronic Systems (ITherm), Orlando, FL, USA, 2023, pp. 1-7

5.2 Chapter 2 References

[1] J. Lee, I. Mudawar, Low-temperature two-phase microchannel cooling for high heat-flux thermal management of defense electronics, IEEE Trans. Components

[2] Tummala RR. Fundamentals of microsystems packaging [M]. McGraw-Hill

[3] Y. Wang et al., "Investigation on Immersion Cooling Solution for Hyper-scale Data Center Application," 2023 22nd IEEE Intersociety Conference on Thermal and Thermomechanical Phenomena in Electronic Systems (ITherm), Orlando, FL, USA, 2023, pp. 1-7, doi: 10.1109/ITherm55368.2023.10177671.I. S. Jacobs and C. P. Bean, “Fine particles, thin films and exchange anisotropy,” in Magnetism, vol. III, G. T. Rado and H. Suhl, Eds. New York: Academic, 1963, pp. 271–350.

[4] Huang R, Masanet E. The uniform methods project: Methods for determining energy efficiency saving for specific measures. Chapter 20: data center IT efficiency measures, Subcontract report NREL/SR-7A40 63181; Jan 2015

[5] ASHRAE Technical Committee 9.9. 2019. Water-Cooled Servers: Common Designs, Components, and Processes

[6] Ogle, M., Curtis, R., “Next-Generation PowerEdge Servers: Thoughtful Thermal Design,” Dell Technologies, Direct from Development Series, 2021

[7] R. Curtis, T. Shedd and E. B. Clark, "Performance Comparison of Five Data Center Server Thermal Management Technologies," 2023 39th Semiconductor Thermal Measurement, Modeling & Management Symposium (SEMI-THERM), San Jose, CA, USA, 2023, pp. 1-9, doi: 10.23919/SEMI-THERM59981.2023.10267908.

[8] Li, Z., & Kandlikar, S., 2015, “Current Status and Future Trends in DataCenter Cooling Technologies,” Heat Transfer Engineering, 36(6), pp. 523-538, DOI: 10.1080/01457632.2014.939032

[9] ASHRAE Technical Committee 9.9, 2021, “Emergence and expansion of liquid cooling in mainstream data centers,” American Society of Heating, refrigeration and Air-Conditioning Engineers, Atlanta, GA.

[10] S. Saini, J. Gullbrand, S. Sarangi, E. McAfee and D. Damm, "Numerical Investigation Of Influence Of Tank Design On Thermal And Flow," 2023 22nd IEEE Intersociety Conference on Thermal and Thermomechanical Phenomena in Electronic Systems (ITherm), Orlando, FL, USA, 2023, pp. 1-7, doi: 10.1109/ITherm55368.2023.10177563.

[11] Saini, S., Shah, J. M., Shahi, P., Bansode, P., Agonafer, D., Singh, P., Schmidt, R., and Kaler, M. (September 15, 2021). "Effects of Gaseous and Particulate Contaminants on Information Technology Equipment Reliability—A Review." ASME. J. Electron. Packag. September 2022; 144(3): 030801. https://doi-org.ezproxy.uta.edu/10.1115/1.4051255

[12] ASHRAE Technical Committee 9.9., 2011, “Particulate and Gaseous Contamination Guidelines for Data Centers,” ASHRAE Inc., Atlanta, GA.

[13] Shah, J. M. , Anand, R. , Saini, S. , Cyriac, R. , Agonafer, D. , Singh, P. , and Kaler, M. , “ Development of a Technique to Measure Deliquescent Relative Humidity of Particulate Contaminants and Determination of the Operating Relative Humidity of a Data Center,” ASME Paper No. IPACK2019-6601.10.1115/IPACK2019-6601

[14] Kheirabadi, Ali & Groulx, Dominic. (2016). Cooling of Server Electronics: A Design Review of Existing Technology. Applied Thermal Engineering. 105.10.1016/j.applthermaleng.20 16.03.056

[15] Xu S, Zhang H, Wang Z. Thermal Management and Energy Consumption in Air, Liquid, and Free Cooling Systems for Data Centers: A Review. Energies. 2023; 16(3):1279. https://doi.org/10.3390/en16031279

[16] Shahi, P., Saini, S., Bansode, P., Agonafer, D., “A comparative study of energy savings in a liquid-cooled server by dynamic control of coolant flow rate at server level,” IEEE Transactions on Components, Packaging and Manufacturing Technology, Vol. 11, No. 4, pp. 616-624, 2021.

[17] Agonafer, D, Bansode, P, Saini, S, Gullbrand, J, & Gupta, A. "Single Phase Immersion Cooling for Hyper Scale Data Centers: Challenges and Opportunities." Proceedings of the ASME 2023 Heat Transfer Summer Conference collocated with the ASME 2023 17th International Conference on Energy Sustainability. ASME 2023 Heat Transfer Summer Conference. Washington, DC, USA. July 10–12, 2023. V001T16A008. ASME. https://doi-org.ezproxy.uta.edu/10.1115/HT2023-107598

[18] Herring, J., Smith, P., Lamotte-Dawaghreh, J., Bansode, P., Saini, S., Bhandari, R., & Agonafer, D. (2022). Machine learning-based heat sink optimization model for single-phase immersion cooling. Proceedings of ASME 2022 International Technical Conference and Exhibition on Packaging and Integration of Electronic and Photonic Microsystems, InterPACK 2022.

[19] Herring, J., Gupta, G., Bansode, P., Lamotte-Dawaghreh, J., & Agonafer, D. (2023). Single-phase immersion cooling multi-design variable heatsink optimization for natural convection. Proceedings of ASME 2023 International Technical Conference and Exhibition on Packaging and Integration of Electronic and Photonic Microsystems, InterPACK 2023. https://doi-org.ezproxy.uta.edu/10.1115/IPACK2023-112019

[20] Gupta, G, Nair, V, Pundla, SA, Bansode, P, Suthar, R, Herring, J, Lamotte-Dawaghreh, J, Sivaraju, KB, Agonafer, D, Mynampati, P, & Sweeney, M. "Optimization of a Air-Cooled Heat sink for Immersion Cooling Application." Proceedings of the ASME 2023 International Technical Conference and Exhibition on Packaging and Integration of Electronic and Photonic Microsystems. ASME 2023 International Technical Conference and Exhibition on Packaging and Integration of Electronic and Photonic Microsystems. San Diego, California, USA. October 24–26, 2023. V001T01A013. ASME. https://doi-org.ezproxy.uta.edu/10.1115/IPACK2023-112054

[21] Saini, S., Wagh, T., Bansode, P., Shahi, P., Herring, J., Lamotte-Dawaghreh, J., Shah, J. M., & Agonafer, D. (2022). A numerical study on multi-objective design optimization of heat sinks for forced and natural convection cooling of immersion-cooled servers. Journal of Enhanced Heat Transfer, 29(8), 15–35.

[22] Chen, F., Jiang, X., Lu, C., Wang, Y., Wen, P., & Shen, Q. (2023). Heat transfer efficiency enhancement of gyroid heat exchanger based on multidimensional gradient structure design. International Communications in Heat and Mass Transfer, 149. https://doi-org.ezproxy.uta.edu/10.1016/j.icheatmasstransfer.2023.107127

[23] Samson, S., Tran, P., & Marzocca, P. (2023). Design and modelling of porous gyroid heatsinks: Influences of cell size, porosity and material variation. Applied Thermal Engineering, 235. https://doi-org.ezproxy.uta.edu/10.1016/j.applthermaleng.2023.121296

[24] Bharadwaj, B, Singh, P, & Mahajan, RL. "Optimal Design of Additively Manufactured Metal Lattice Heat Sinks for Electronics Cooling." Proceedings of the ASME 2022 Heat Transfer Summer Conference collocated with the ASME 2022 16th International Conference on Energy Sustainability. ASME 2022 Heat Transfer Summer Conference. Philadelphia, Pennsylvania, USA. July 11–13, 2022.

[25] Sajjad, U., Rehman, T., Ali, M., Park, C. W., & Yan, W.-M. (2022). Manufacturing and potential applications of lattice structures in thermal systems: A comprehensive review of recent advances. International Journal of Heat and Mass Transfer, 198. https://doi-org.ezproxy.uta.edu/10.1016/j.ijheatmasstransfer.2022.123352

[26] Ning, J., Wang, X., Huang, H., Wang, S., & Yan, W. (2023). Topology optimized novel additively manufactured heat sink: Experiments and numerical simulations. Energy Conversion and Management, 117024. https://doi-org.ezproxy.uta.edu/10.1016/j.enconman.2023.117024

[27] See, Y. S., Ho, J. Y., Leong, K. C., & Wong, T. N. (2022). Experimental investigation of a topology-optimized phase change heat sink optimized for natural convection. Applied Energy, 314. https://doi-org.ezproxy.uta.edu/10.1016/j.apenergy.2022.118984

[28] M. Frank, M. Matthews, J. Madril and I. Winfield, "ELECTROCHEMICAL ADDITIVE MANUFACTURING: A NOVEL APPROACH TO THERMAL MANAGEMENT OF ELECTRONICS," 2023 IEEE 73rd Electronic Components and Technology Conference (ECTC), Orlando, FL, USA, 2023, pp. 432-436, doi: 10.1109/ECTC51909.2023.00078.

[29] Pope, Daniel., Miyoshi, Mark “Single-Phase Immersion Cooling: The Path to 1000W TDP & Beyond”, Case Study – Submer, https://submer.com/wp-content/uploads/2023/10/Intel-and-Submer-FCHS-Case-Study.pdf

[30] ©2024 IEEE. Reprinted, with permission, from Joseph Herring, Jacob Lamotte-Dawaghreh, Gautam Gupta, Dereje Agonafer, Joseph Madril, Tim Ouradnik, Ian Winfield, Michael Matthews, CFD Evaluation of Electrochemical Additively Manufactured Heat Sinks for Single-Phase Immersion Cooling, 2024 23rd IEEE Intersociety Conference on Thermal and Thermomechanical Phenomena in Electronic Systems (ITherm), May 2024

5.3 Chapter 3 References

[1] Shah, J. M., Dandamudi, R., Bhatt, C., Rachamreddy, P., Bansode, P., & Agonafer, D. (2019). CFD analysis of thermal shadowing and optimization of heatsinks in 3rd generation open compute server for single-phase immersion cooling. ASME 2019 International Technical Conference and Exhibition on Packaging and Integration of Electronic and Photonic Microsystems, InterPACK 2019.

[2] Sarangi, S., McAfee, E. D., Damm, D. G., & Gullbrand, J. (2022). Single-Phase Immersion Cooling Performance in Intel Servers with Immersion Influenced Heatsink Design. 38th Annual Semiconductor Thermal Measurement, Modeling and Management Symposium, SEMI-THERM 2022 - Proceedings, 1–5.

[3] Vuckovic, M., & Depret, N. (2016). Impacts of local cooling technologies on air cooled data center server performance: test data analysis of Heatsink, Direct Liquid Cooling and passive 2-Phase Enhanced Air Cooling based on Loop Heat Pipe. 2016 32nd Thermal Measurement, Modeling & Management Symposium (SEMI-THERM). Proceedings, 71–80.

[4] Shinde, P. A., Bansode, P. V., Saini, S., Kasukurthy, R., Chauhan, T., Shah, J. M., & Agonafer, D. (2019). Experimental analysis for optimization of thermal performance of a server in single phase immersion cooling. ASME 2019 International Technical Conference and Exhibition on Packaging and Integration of Electronic and Photonic Microsystems, InterPACK 2019.

[5] Saini, S., Wagh, T., Bansode, P., Shahi, P., Herring, J., Lamotte-Dawaghreh, J., Shah, J. M., & Agonafer, D. (2022). A NUMERICAL STUDY ON MULTI-OBJECTIVE DESIGN OPTIMIZATION OF HEAT SINKS FOR FORCED AND NATURAL CONVECTION COOLING OF IMMERSION-COOLED SERVERS. Journal of Enhanced Heat Transfer, 29(8), 15–35.

[6] Herring, J., Smith, P., Lamotte-Dawaghreh, J., Bansode, P., Saini, S., Bhandari, R., & Agonafer, D. (2022). MACHINE LEARNING-BASED HEAT SINK OPTIMIZATION MODEL for SINGLE-PHASE IMMERSION COOLING. Proceedings of ASME 2022 International Technical Conference and Exhibition on Packaging and Integration of Electronic and Photonic Microsystems, InterPACK 2022.

[7] Yazicioglu, B., & Yuncu, H. (2007). Optimum fin spacing of rectangular fins on a vertical base in free convection heat transfer. Heat and Mass Transfer, 44(1), 11–21.

[8] Jang, D., Yu, S.H., and Lee, K.S., Multidisciplinary Optimization of a Pin-Fin Radial Heat Sink for LED Lighting Applications, Int. J. Heat Mass Transf., vol. 55, no. 4, pp. 515–521, 2012.

[9] Will, J. and Most, T., Metamodel of Optimized Prognosis (MoP)—An Automatic Approach for User Friendly Parameter Optimization, in 6th Optimization and Stochastic Days, Weimar, Germany, October 15–16, 2009.

[10] Murthy, P., Gupta, G., Herring, J., Lamotte-Dawaghreh, J., Sivaraju, K. B., Bansode, P., Modi, H., Agonafer, D., Mynampati, P., & Sweeney, M. (2022). CFD SIMULATION-BASED COMPARATIVE STUDY of FORCED CONVECTION SINGLEPHASE LIQUID IMMERSION COOLING for A HIGH-POWERED SERVER. Proceedings of ASME 2022 International Technical Conference and Exhibition on Packaging and Integration of Electronic and Photonic Microsystems, InterPACK 2022.

[11] Sivaraju, K. B., Bansode, P., Gupta, G., Lamotte-Dawaghreh, J., Saini, S., Simon, V., Herring, J., Karajgikar, S., Mulay, V., & Agonafer, D. (2022). COMPARATIVE STUDY of SINGLE-PHASE IMMERSION COOLED TWO SOCKET SERVER in TANK and SLED CONFIGURATIONS. Proceedings of ASME 2022 International Technical Conference and Exhibition on Packaging and Integration of Electronic and Photonic Microsystems, InterPACK 2022.

[12] Herring, J., et al. (2023). “Single-Phase Immersion Cooling Multi-Design Variable Heat Sink Optimization for Natural Convection” ASME 2023 International Technical Conference and Exhibition on Packaging and Integration of Electronic and Photonic Microsystems

5.4 Chapter 4 References

[1] Miller, R. (2022). After Record Growth in 2021, Data Centers Are Building Bigger for 2022. Data Center Frontier, https://datacenterfrontier.com/after-record-growth-in-2021-data-centers-are-building-bigger-for 2022/#:~:text=The%20installed%20base%20of%20hyperscale,will%20continue%20growing%20rapidly%20thereafter.

[2] Bawden, T. (2016). Global warming: Data centres to consume three times as much energy in next decade, experts warn. INDEPENDENT. https://www.independent.co.uk/climate-change/news/global-warming-data-centres-to-consume-three-times-as-much-energy-in-next-decade-experts-warn-a6830086.html

[3] Greenberg, S., Mills, E., Tschudi, B., Rumsey, P. and Myatt, B., Best practices for data centers: Lessons learned from benchmarking 22 data centers,*Proceedings of the ACEEE Summer Study on Energy Efficiency in Buildings in Asilomar, CA. ACEEE, August*,*3*,pp.76-87, 2006

[4] Wen, F., 2018, “Best Practice of Alibaba Datacenter-Immersion Cooling Escorts Cloud Computing,” Presented at the OCP Global Summit, San Jose, CA, p. 12, accessed Sept. 3, 2021, https://www.opencompute.org/files/Immersion-Cooling-for-Green-Computing-V1.0.pdf

[5] Shahi, P., Saini, S., Bansode, P., Agonafer, D., A Comparative Study of Energy Savings in a Liquid-Cooled Server by Dynamic Control of Coolant Flow Rate at Server Level, in*IEEE Transactions on Components, Packaging and Manufacturing Technology*, vol. 11, no. 4, pp. 616-624, doi:10.1109/TCPMT.2021.3067045, 2021

[6] Shahi, P., Deshmukh, A. P., Hurnekar, H. Y., Saini, S., Bansode, P., Kasukurthy, R., and Agonafer, D., Design, Development, and Characterization of a Flow Control Device for Dynamic Cooling of Liquid-Cooled Servers, ASME.*J. Electron. Packag*., **vol** 144(4): 041008.https://doi.org/10.1115/1.4052324, 2022

[7] Niazmand, A., Chauhan, T., Saini, S., Shahi, P., Bansode, P.V., & Agonafer, D., CFD Simulation of Two-Phase Immersion Cooling Using FC-72 Dielectric Fluid, Proceedings of the ASME 2020 International Technical Conference and Exhibition on Packaging and Integration of Electronic and Photonic Microsystems. ASME 2020 International Technical Conference and Exhibition on Packaging and Integration of Electronic and Photonic Microsystems. Virtual, Online. October 27–29, 2020. V001T07A009. ASME. https://doi.org/10.1115/IPACK2020-2595

[8] Hoang, C.H., Hoang, C.H., Tradat, M., Manaserh, Y., Ramakrisnan, B., Rangarajan, S., Hadad, Y., Schiffres, S. and Sammakia, B., A Review of Recent Developments in Pumped Two-Phase Cooling Technologies for Electronic Devices," in IEEE Transactions on Components,

Packaging and Manufacturing Technology, vol. 11, no. 10, pp. 1565-1582, doi: 10.1109/TCPMT.2021.3117572, 2021

[9] Yazicioğlu, B. and Yüncü, H., 2007. Optimum fin spacing of rectangular fins on a vertical base in free convection heat transfer.*Heat and Mass Transfer*,*44*(1), pp.11-21, https://doi.org/10.1007/s00231-006-0207-6, 2007

[10] Jang, D., Yu, S.H. and Lee, K.S., 2012. Multidisciplinary optimization of a pin-fin radial heat sink for LED lighting applications.*International Journal of Heat and Mass Transfer*,*55*(4), pp. 515-521, https://doi.org/10.1016/j.ijheatmasstransfer.2011.11.016, 2012

[11] Kim, D.K., Thermal optimization of plate-fin heat sinks with fins of variable thickness under natural convection,*International journal of heat and mass transfer*,*55*(4), pp.752-761, https://doi.org/10.1016/j.ijheatmasstransfer.2011.10.034, 2012

[12] Subasi, A., Sahin, B. and Kaymaz, I., Multi-objective optimization of a honeycomb heat sink using Response Surface Method.*International Journal of Heat and Mass Transfer*,*101*, pp.295-302, https://doi.org/10.1016/j.ijheatmasstransfer.2016.05.012, 2016

[13] Max Mowbray, Thomas Savage, Machine learning for biochemical engineering: A review, *Biochemical Engineering Journal 172 (2021) 108054.*

[14] Shashank Reddy Vadyala, Sai Nethra Betger, A review of physics-based machine learning in civil engineering, *Results in Engineering 13 (2022)100316*

[15] Rakshit, S., Clement, N., Vajjhala, N.R. (2022). Exploratory Review of Applications of Machine Learning in Finance Sector. In: Borah, S., Mishra, S.K., Mishra, B.K., Balas, V.E., Polkowski, Z. (eds) Advances in Data Science and Management . Lecture Notes on Data Engineering and Communications Technologies, vol 86. Springer, Singapore. https://doi-org.ezproxy.uta.edu/10.1007/978-981-16-5685-9_12

[16] Javed Azmi, Muhammad Arif, Md Tabrez Nafis, M. Afshar Alam, Safdar Tanweer, Guojun Wang, A systematic review on machine learning approaches for cardiovascular disease prediction using medical big data, *Medical Engineering and Physics 105 (2022) 103825*

[17] Anastasiia Grishina, Marta Chinnici, Ah-Lian Kor, Eric Rondeau, and Jean-Philippe Georges, A Machine Learning Solution for Data Center Thermal Characteristics Analysis,*Energies (2020), 13, 4378*

[18] Qiu Fang,Zhe Li,Yaonan Wang,Mengxuan Song, Jun Wang, A neural-network enhanced modeling method for real-time evaluation of the temperature distribution in a data center, *Neural Computing and Applications, (2019) 31:8379-8391*

[19] Zeeshan Rasheed, Wei Xiong, Gaoxiang Cong, Hongxun Niu, Jianxiong Wan, Yongsheng Wang, and Lixiao Li, Performance Evaluation and Machine Learning based Thermal Modeling of Tilted Active Tiles in Data Centers, *ACM International Conference Proceeding Series, p 32-38, June 19, 2020, ICMLT 2020*

[20] Shashikant Ilager, Kotagiri Ramamohanarao, and Rajkumar Buyya , Fellow, IEEE, Thermal Prediction for Efficient EnergyManagement of Clouds Using Machine Learning, *IEEE TRANSACTIONS ON PARALLEL AND DISTRIBUTED SYSTEMS, VOL. 32, NO. 5, MAY 2021*

[21] Masayuki Nakamura, Learning and Optimization Models for Energy Efficient Cooling Control in Data Center,*2016 55th Annual Conference of the Society of Instrument and Control Engineers of Japan (SICE), p 395-400, 2016*

[22] Peng Xiao1, Zhenyu Ni, Dongbo Liu, Zhigang Hu, A power and thermal-aware virtual machine management frameworkbasedon machine learning, *Cluster Computing, v 24, n 3, p 2231-2248, September 2021*

[23] SeyedMorteza, MirhoseiniNejad, Ghada Badawy, Douglas G. Down, Holistic thermal-aware workload management and infrastructurecontrol for heterogeneous data centers usingmachine learning, *Future Generation Computer Systems, v 118, p 208-218, May2021*

[24] Vibin Shalom Simon, Ashwin Siddarth, Dereje Agonafer, Artificial Neural Network Based Prediction of Control Strategies for Multiple Air-Cooling Units ina Raised-floor Data Center, *2020 19th IEEE Intersociety ConferenceonThermalandThermomechanical Phenomena in Electronic Systems (ITherm), p 334-40, 2020*

[25] Acharya, P.V. (University of Texas at Austin, Walker Department of Mechanical Engineering, Austin, TX, United States); Lokanathan, M.; Ouroua, A.; Hebner, R.; Strank, S.; Bahadur, V. Source: 2021 20th IEEE Intersociety Conference on Thermal and Thermomechanical Phenomena in Electronic Systems (iTherm), p 162-71,2021

[26] Acharya, P.V. (University of Texas at Austin, Walker Department of Mechanical Engineering, Austin, TX TX 7871, United States); Lokanathan, M.; Ouroua, A.; Hebner, R.; Strank, S.; Bahadur, V. Source: Journal of Electronic Packaging, v 143, n 4, p041109 (14 pp.), 2021

[27] McWilliams, T.D., Evaluating Heat Sink Performance in an Immersion-Cooled Server System, M.S Thesis, The University of Texas at Arlington, USA, 2014

[28] Martın Abadi, Ashish Agarwal, Paul Barham, Eugene Brevdo, Zhifeng Chen, Craig Citro, Greg S. Corrado, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Ian Goodfellow, Andrew Harp, Geoffrey Irving, Michael Isard, Yangqing Jia, Rafal Jozefowicz, Lukasz Kaiser, Manjunath Kudlur, Josh Levenberg, Dandelion Mane, Rajat Monga, Sherry Moore, Derek Murray, Chris Olah, Mike Schuster, Jonathon Shlens, Benoit Steiner, Ilya Sutskever, Kunal Talwar, Paul Tucker, Vincent Vanhoucke, Vijay Vasudevan, Fernanda Viegas, Oriol Vinyals, Pete Warden, Martin Wattenberg, Martin Wicke, Yuan Yu, and Xiaoqiang Zheng. TensorFlow: Large-scale machine learning on heterogeneous systems, 2015. URL https://www. tensorflow.org/. *Software available from tensorflow.org*

[29] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Re- search, 12:2825–2830, 2011*

[30] Herring, J., et al. (2022). “Machine Learning-Based Heat Sink Optimization Model for Single-Phase Immersion Cooling” ASME 2022 International Technical Conference and Exhibition on Packaging and Integration of Electronic and Photonic Microsystems

Biography

My name is Joseph Herring. In my full-time role as Senior Hardware Design Engineer at Lockheed Martin Missiles and Fire Control, I support Applied Research and Advanced Programs Mechanical Design, Electromechanical Packaging, and Thermal Hardware Design/Analysis.

My Ph.D. research at The University of Texas at Arlington includes high power-density electronics cooling, data center thermal management, heat sink optimization, air cooling, single-phase liquid immersion cooling, direct-to-chip liquid cooling, and machine learning optimization modeling.