

See discussions, stats, and author profiles for this publication at: <https://www.researchgate.net/publication/361234166>

A Numerical Study on Multi-objective Design Optimization of Heatsinks for Forced and Natural Convection Cooling of Immersion Cooled Servers

Article in *Journal of Enhanced Heat Transfer* · January 2022

DOI: 10.1615/JEnhHeatTransf.2022043805

CITATION

1

READS

41

8 authors, including:



Satyam Saini

Intel

23 PUBLICATIONS 109 CITATIONS

[SEE PROFILE](#)



Pratik Bansode

University of Texas at Arlington

20 PUBLICATIONS 94 CITATIONS

[SEE PROFILE](#)



Pardeep Shahi

University of Texas at Arlington

16 PUBLICATIONS 72 CITATIONS

[SEE PROFILE](#)



Jimil M. Shah

TMGcore

38 PUBLICATIONS 224 CITATIONS

[SEE PROFILE](#)

Some of the authors of this publication are also working on these related projects:



Reliability of Electronic Packages [View project](#)



Dynamic cold plate design for Chip Scale Packages [View project](#)

A NUMERICAL STUDY ON MULTI-OBJECTIVE DESIGN OPTIMIZATION OF HEAT SINKS FOR FORCED AND NATURAL CONVECTION COOLING OF IMMERSION-COOLED SERVERS

Satyam Saini,^{1,*} Tushar Wagh,¹ Pratik Bansode,¹
Pardeep Shahi,¹ Joseph Herring,¹ Jacob Lamotte-Dawaghreh,¹
Jimil M. Shah,² & Dereje Agonafer¹

¹Mechanical and Aerospace Engineering Department, The University of Texas at
Arlington, 500 W 1st St., Arlington, Texas 76019, USA

²TMGCore Inc., Plano, Texas 75024, USA

*Address all correspondence to: Satyam Saini, Mechanical and Aerospace Engineering
Department, The University of Texas at Arlington, 500 W 1st St., Arlington, TX 76019,
USA, E-mail: satyam.saini@mavs.uta.edu

Original Manuscript Submitted: 3/27/2022; Final Draft Received: 5/23/2022

Rapidly rising computational, storage and networking demands have brought about an increment in both the number and energy density of modern data centers. Typical air-cooled high-performance servers require low air-supply temperatures as well as higher air-flow rates making air-cooling inefficient above certain thermal design power values. Single-phase immersion cooling solves this issue by offering significantly higher thermal mass and known reliability enhancements. When choosing to implement immersion cooling for server hardware that is designed for air-cooling, immersion-specific optimized heat sinks should be used. This study presents an in-depth investigation of multi-objective and multi-design variable optimization of heat sinks for immersion-cooled servers for a fixed pumping power using computational fluid dynamics. The optimization is conducted for both copper and aluminum heat sinks by using pressure drop and thermal resistance minimization as objective functions. Differences in the optimized values of the heat-sink geometry were compared and quantified for natural as well as forced convection cooling. The results show that the total effect values of heat-sink geometric parameters vary significantly between forced and natural convection optimized heat sinks. Approximately 15% improvement in both thermal resistance and pressure drop values was observed in the optimized heat sink. For the case of natural convection cooling, it is observed that the heat-sink height has a greater significance than the fin thickness.

KEY WORDS: heat sink, natural convection, thermal management, servers, immersion cooling

1. INTRODUCTION

Data centers have quickly become the backbone of any modern economy with the emergence of technologies like cloud computing, online media, social networking, artificial intelligence (AI), and machine learning. Active internet user statistics suggest that at any given time, around 5 billion users are active on the internet, which sheds ample light on the reliance on network servers and data centers (Internet Live Stats, 2022). Data center energy trends, while having slowed down due to efficient IT equipment (Masanet et al., 2020), continue to rise every year. It is however uncertain if the cooling and energy efficiency enhancements can continue to offset the rapid growth in computing workloads (Shehabi et al., 2018). Power densities in a typical data center can be 15–100 times greater than in generic commercial buildings (Greenberg et al., 2006). The growth in this demand has also environmental implications like an increase in greenhouse emissions and excessive water usage, both direct and indirect (Siddik et al., 2021).

In the last decade, as the energy demands and processor power densities have increased due to higher processing demands, traditional air-cooling methods have been limited to cooling processors with lower thermal design power (TDP) limits. With rising energy consumption and increasing complexity in thermal management, researchers have proposed various strategies to enhance thermal performance. Some of them include using nanofluids in immersion and cold-plate based cooling (Shahi et al., 2020; Niazmand et al., 2020a), increasing power savings using different operation strategies of current cooling techniques (Shahi et al., 2021, 2022), or using liquid-based phase-change cooling technologies (Niazmand et al., 2020b; Hoang et al., 2021). Among the popular liquid-cooling technologies, single-phase immersion cooling stands out owing to its ease of deployment, low costs of dielectric fluids (as compared to two-phase immersion-cooling fluids), and low-complexity cooling infrastructure (Shah et al., 2016). Many of the present proprietary cooling solutions have established the efficacy of single-phase immersion cooling in terms of low-power usage effectiveness (Eiland et al., 2014; Shinde et al., 2019; Gandhi et al., 2019) values and enhanced reliability of server components (Shah et al., 2022). The ability of the single-phase immersion cooling in extreme temperatures (Bansode et al., 2018, 2020) and also a reduction of server form factor (Shah et al., 2019) has also been studied extensively. A key direct advantage of complete submersion of servers in dielectric fluids is that it disconnects the server components from the harsh environment such as gaseous contaminants (Saini et al., 2022), reduces failures due to fan vibrations, and removes the necessity of cooling peripheral components as hot components are in direct contact with the coolant. Immersion cooling offers significant advantages as compared to air-cooling but requires careful thermal and non-thermal design considerations for air-cooled hardware (Brink et al., 2022; Shah et al., 2019). As an example, when immersing air-cooled hardware, fans should be removed from the server, hard drives need to be sealed, material compatibility issues should be addressed, and the heat sink design should be optimized, which is the objective of this study.

To achieve consistently reliable and peak performance from the central processing unit (CPU) or graphics processing unit (GPU) in immersion cooling, an optimized heat sink needs to be utilized, rather than using an air-cooled heat sink. Design optimizations for parallel plate-fin heat sinks have been widely studied including optimization of geometric properties (Yazicioğlu et al., 2007; Jang et al., 2012) and also from a single or multi-objective optimization point of view (Kim, 2012; Subasi et al., 2016). Optimum geometric parameters such as fin spacing, fin thickness, base thickness, and fin count play an important role in maximizing the thermal performance of the heat sink as well as the processors and saving the pumping power by reducing

the pressure drop across the heat sink. Current literature shows an increasing number of studies related to the optimization of heat-sink solutions for both air- and liquid-cooled systems using various computational fluid dynamics (CFD)-based and analytical methods.

Chen and Chen used a multi-objective and novel direction-based algorithm to optimize plate-fin heat sinks integrated with an impingement fan using a commercially available multiphysics tool (Chen and Chen, 2013). The optimized parallel plate heat sink showed both superior heat-transfer performance and reduced weight. Fuzzy logic-based approaches have also been used to quantify the effect of heat-sink design parameters on thermal performance (Chiang et al., 2006). Design parameters of a pin-fin heat sink such as fin spacing, pin-fin diameter, and height were investigated experimentally. Analysis of variance (ANOVA) was then used to explore the effect of these design parameters on heat-sink characteristics like thermal resistance, pressure drop, and average heat-transfer coefficient. Response surface method (RSM) was used by Chiang and Chang (2006) to obtain optimum design parameters for pin-fin heat sinks to achieve higher thermal performance. Minimizing entropy generation rate as an objective function was done by Chen et al. (2008). They optimized a plate-fin heat sink for a CPU using a coded genetic algorithm for obtaining optimum design parameters for the heat sink. In terms of objective functions, Devi et al. (2012) used a Taguchi-based non-gradient method for minimizing three objective functions, namely, radiation emission, thermal resistance, and heat-sink mass. Single-objective, multi-parametric study of thermal optimization at the server level has also been conducted (Shah et al., 2019). This study looked at different heat-sink fin designs and parametrized the heat-sink design variables to reduce only the overall thermal resistance of the heat sink. However, the pressure drop across the heat sink or hydraulic performance of the heat sink was not used as an objective function.

However, these studies are very generic in terms of the application of the heat sinks and, to the best of the authors' knowledge, no study in the literature discusses the heat-sink optimization specific to immersion-cooled servers for both natural and forced convection cooling. The main purpose of this investigation is to dive deep into different design parameters and objective function combinations for optimizing a parallel plate-fin heat sink of an air-cooled Open Compute server that is used for immersion cooling. A CFD model of the server was developed using ANSYS Icepak and was validated against the experimental data in a previously published work for thermal performance. Baseline CFD simulations were first done using the air-cooled heat sink in the dielectric fluid, EC-100 (Engineered Fluids), for both natural and forced convection cooling modes. The heat-sink material was also varied from aluminum to copper in the baseline simulations for both the cooling modes. OptiSLang, a dedicated design optimization and analysis tool inside the ANSYS platform, was used for heat-sink design optimization. Heat-sink fin thickness, fin count, fin height, and base thickness were optimized in different combinations at constant pumping power. The objective functions, heat-sink pressure drop, and thermal resistance are defined in Icepak and set to be minimized for an optimized design. A range of values of the heat-sink design parameters is defined in Icepak which are exported to OptiSLang along with the information related to objective functions and design parameters. OptiSLang then creates a design of experiment (DOE) based on the input range of the design variables using a full-factorial approach to create multiple design points. These design points are then solved iteratively and the response surfaces, effect plots, and Pareto fronts are generated by OptiSLang based on the results from each of the design points. Analysis of the effect of each of the design variables on the objective functions is done to establish which parameters have a greater impact on the performance of the optimized copper and aluminum heat sinks under natural and forced convection.

2. NUMERICAL MODELING SETUP

2.1 CFD Model Validation

A computational model of an Open Compute server was designed in ANSYS Icepak by simplifying the server and considering components critical to the heat-transfer process like CPU, heat sinks, and dual in-line memory modules (DIMMs) as shown in Fig. 1. Each CPU unit in the CFD model is designed to dissipate 115 W. The thermal stack representing the CPU consists of a 2D heat source at the bottom of a heat sink with a thermal interface material (TIM) of a thickness of 0.2 mm with a thermal conductivity value of 3.8 W/mK. The CFD model was first validated against an experimental study done by our research group on the same server (McWilliams, 2014). Figure 2 shows the difference in the values of the maximum temperature at different flow rates between the benchmark experimental study and the CFD simulation data. A maximum variation of 2% was observed in the maximum temperature value between the CFD and the experimental study. Some differences seen in the results could be because the heat sinks in the server used during the experiments have embedded heat pipes in them. However, heat




FIG. 1: (Top) Computational model of the server showing the heat sinks and the memory modules and (bottom left) air-cooled version of the real server; (bottom right) front view of the computational model showing the heat-sink fins and DIMMs




FIG. 2: A comparison of the average junction temperatures (maximum source temperatures in case of CFD model) between the experimental and simulation data

pipes were not modeled in the CFD study to keep the heat-sink design simpler and accelerate the optimization process. Another difference is that the baseline study used mineral oil as the cooling fluid but a synthetic dielectric fluid EC-100 (higher heat capacity than mineral oil) was used for the present study. Therefore, there was not a significantly large change in source temperatures and the temperature variation also shows the same trend as the experimental data.

A grid independence study was also done for the baseline CFD model to determine the validity and precision of the results obtained. Coarse to fine-level meshes were generated by reducing the minimum element size in each of the three directions. ANSYS Icepak, by default, generates a minimum element size of 1/20 of the length in that specific direction. For the grid independence study, this element size was reduced to 1/5 of the length for a coarse mesh. For a refined grid, minimum mesh element sizes of 1/30, 1/40, 1/50, and 1/60 of the length were generated. To verify the grid independence, both pressure drop across the server and thermal resistance of the heat sink were monitored for different mesh element count as shown in Fig. 3. It can be seen that both the pressure drop and thermal resistance show a non-varying trend at all the mesh counts. Therefore, the default meshing size of 1/20 was used for the simulations.

2.2 OptiSLang Setup

OptiSLang is used as the design optimization tool in the present investigation. As a part of ANSYS Workbench, OptiSLang has the advantage that it can be used for direct integration with any of the ANSYS thermal, structure, electrical, or fluid tools. The simulation model to be optimized is solved independently in any of the ANSYS tools or modules. Design parameters and their range or bounds for optimization are also defined in the simulation module itself and are then imported to OptiSLang. The software uses a meta-modeling approach for sampling the design space known as the adaptive meta-model of optimal prognosis (AMOP), which uses a coefficient of prognosis (CoP) for the quality approximation of the model (Mimery, 2020). It can be mathematically represented as

$$\text{CoP} = 1 - \frac{SS_E^{\text{Prediction}}}{SS_T} \quad (1)$$




FIG. 3: Variation of the pressure drop and thermal resistance values with changing mesh element count

Here, SS_T is equivalent to total variation and SS_E represents variation due to regression or the sum of the square of prediction errors. The higher the value of CoP the more accurate is the model's representation of the data. This reduces the post-processing of the output data from the design exploration space and helps in the direct assessment of the response surface models (Will and Most, 2009). Also, typical meta-model approaches provide limited accuracy when the number of input variables starts increasing, thereby making the design sample space increase exponentially. This can be overcome by using AMOP, which improves the prediction quality of an approximation model by eliminating the design variables that are not important in the model (Most and Will, 2011). A summary of the simulation model setup and its integration with the design optimization tool is shown in detail using a schematic in Fig. 4.




FIG. 4: Overview of the integration of CFD simulation model setup and optimization setup in OptiSlang

The objective functions and the constraints that control the optimization process can be defined in the simulation module as well as OptiSLang. Once the optimization process is completed, a sensitivity analysis is done to determine the effect of the chosen design variables on the objective functions. As the baseline design in the present investigation is fixed a pre-determined upper bound of the design variables was selected. Using a very large range of design variables also makes it difficult to assess the design exploration space. The optimized design points are visualized on a Pareto front for each of the optimizations run to determine the design points that offer the least thermal resistance and pressure drop.

3. METHODOLOGY

The primary purpose of this study is to investigate various possible design optimization methodologies for heat sinks when transitioning from air-cooling to immersion cooling. In this section, the underlying methodology, assumptions, and different parametric cases run for the optimized heat sink are discussed. The optimization of the heat sinks was carried out for a constant flow rate of 2 lpm and peak utilization power for both the CPUs at 115 W. The baseline air-cooled heat sink has a fin count of 41, a base thickness of 4 mm, and 0.23 mm thick fins with a fin height of 37 mm. The dielectric fluid selected for the CFD study is a commercially available synthetic fluid, EC100. The optimization was done for both copper and aluminum heat sinks under natural and forced convection flow regimes separately. This was done to quantify the differences in the optimized heat-sink design and the effect of design variables on the objective functions when either the flow regime or heat-sink material is changed. The results of doing these permutations of heat-sink material under different cooling modes will also allow users to directly choose the design variables that have the most impact on the objective functions.

The CFD tool used solves Navier–Stokes equations of mass, momentum, species, and energy to calculate heat transfer in laminar flow conditions. Additional transport equations of turbulence and radiation can be used if the flow and heat transfer involve these phenomena, which was not the case for the current study. These equations are written as follows:

Mass conservation:

$$\frac{\partial \rho}{\partial t} + \nabla \cdot (\rho \vec{v}) = 0 \quad (2)$$

The above equation reduces to $\nabla \cdot (\vec{v}) = 0$ for incompressible fluids.

Momentum equation:

$$\frac{\partial}{\partial t} (\rho \vec{v}) + \nabla \cdot (\rho \vec{v} \vec{v}) = -\nabla p + \nabla \cdot (\bar{\tau}) + \rho \vec{g} + \vec{F} \quad (3)$$

Energy equation:

$$\frac{\partial}{\partial t} (\rho h) + \nabla \cdot (\rho h \vec{v}) = \nabla \cdot [(k + k_t) \nabla T] + S_h \quad (4)$$

Here, the fluid energy equation is written in terms of sensible enthalpy, h . k is the molecular conductivity and k_t is the turbulence transport conductivity. The source term S_h represents user-defined volumetric heat sources. For the solid regions, the energy equation due to conduction within the solid looks as follows:

$$\frac{\partial}{\partial t} (\rho h) = \nabla \cdot (k \nabla T) + S_h \quad (5)$$

Here, k is the thermal conductivity of the solid, ρ is the density, T is the temperature, and S_h is the source term for volumetric heat sources.

Buoyancy-driven flows in mixed or fully natural convection-driven flows are simulated using the Boussinesq model for natural convection. This model treats density as constant in all solved equations except the buoyancy term in the momentum equation:

$$(\rho - \rho_0) g \approx -\rho_0 \beta (T - T_0) g \quad (6)$$

In the above equation, ρ_0 is the constant density of the fluid, T_0 is the operating temperature, and β is the volume expansion coefficient of the fluid.

Figure 5 shows the boundary conditions used for thermal simulations in ANSYS Icepak. The gravity is set in the negative z direction, representing the vertical orientation of the server, as in the case of typical immersion-cooled tanks. The first step in the design optimization process was validating the CFD data with the experimental results as discussed in Section 2.1. After the baseline simulations for the air-cooled heat sink, the bounds for the design variables and the objective functions are defined. Thermal resistance and pressure drop are the typical performance metrics that define the heat-sink performance. Both these functions were defined to be minimized during the optimization process. A summary of the bounds of the design variables, the objective functions, and other variable parameters used in this study is shown in Tables 1 and 2. The




FIG. 5: (Top left) Boundary conditions used for the CFD simulations in ANSYS Icepak and (top right) integration of OptiSLang module with Icepak in ANSYS Workbench; (bottom) boundary conditions used for the natural convection model

TABLE 1: Variable input parameters used in ANSYS Icepak

No.	Parameters	Factor
1	Inlet fluid velocity	Varying cabinet height
2	Heat-sink overall height	Optimization design variable
3	Fin thickness	Optimization design variable
4	No. of fins	Optimization design variable
5	Mesh	Varying cabinet height
6	Thermal resistance	Objective function
7	Pressure drop	Objective function

TABLE 2: Inputs of design variables used for design exploration in OptiSLang (bold and italics are baseline values of the parameters)

S. No.	Overall heat-sink height (mm)	Fin thickness (mm)	No. of fins
1	26	<i>0.23</i>	25
2	29	0.32	27
3	32	0.41	29
4	35	0.5	31
5	38	0.59	33
6	<i>41</i>	0.68	<i>35</i>
Step size	3	0.09	2
Discrete values	6	6	6
Total number of design points $(6 \times 6 \times 6) = 216$			

final CFD results from thermal simulations along with the design variable bounds and objective functions are published to OptiSLang where the multi-objective optimization is performed.

Once the input parameters or the design variables are read in OptiSLang, a design of experiments (DoE) is created. Based on this DoE, a sensitivity analysis is first done to determine the effect of the chosen design variables on the objective functions of the study. This phase is also known as the design exploration phase where the solution space can be sampled using the available sampling methods in OptiSLang. For the current study, AMOP is used for this purpose, the details of which are discussed in Section 2.2. Figure 6 shows the criteria set up in OptiSLang showing the input design variable parameters and the objective function. The outputs from this stage are the response surfaces, response plots, and the total effects matrix. These results help in determining the relationships and the weightage of those relationships between the design variables and the objective functions. The final result of the optimization of the design space is directly obtained on a Pareto front with its axis as the two objective functions. The boundaries of the Pareto front represent the minimum values of the objective functions.

4. RESULTS

The optimization results presented in this section are divided into two subsections: the results for aluminum heat sink and copper heat sink. The results for both the heat sinks also include the




FIG. 6: Setup of optimization parameters in OptiSLang using AMOP for design exploration

results for natural and forced convection. Assessment of the optimized heat-sink design is done based on analyzing surface response plots, CoP, and by compiling the complete design space on a Pareto front. The criteria for optimization of heat sinks and design variables were kept the same for both natural and forced convection heat sinks so that a comparison can be made under similar constraints. As discussed in Section 3, 216 design points were generated for each simulation case with a total of 6 different simulation cases.

4.1 Optimization for Aluminum Heat Sink

The first part of the optimization study was the sensitivity analysis of the design variables to the objective functions. Figure 7 shows the total effects plots for the aluminum heat-sink optimization case under forced convection flow. The effect plot essentially quantifies the impact of each of the inputs on the outputs or post-processing functions. It can be seen that thermal resistance is primarily dependent on heat-sink fin thickness. Similarly, in the case of pressure drop, heat-sink height is a major factor. The fin count of the heat sink does not have a significant impact on both the pressure drop and thermal resistance. This was against the assumption that a large number of fins and a correspondingly larger surface area will improve heat transfer. This is also a part of the reason why heat sinks designed for air-cooling are typically not able to provide the expected heat dissipation in immersion cooling. It can be seen that a linear regression CoP value of




FIG. 7: Total effects plot for aluminum heat sink under forced convection

greater than 99.6% is achieved for all model outputs. This indicates that the sample points were generated based on the design variable inputs and a highly robust model was created.

The 2D and 3D dependencies of the input design variables are visualized using linear regression-based plots and response surfaces. Figure 8 shows the dependency of the input variables on thermal resistance. As observed from the total effects matrix fin count has a very low impact on thermal resistance variation and tends to become asymptotic after a fin count of 33 fins in the current heat-sink design. Discrete values of fin height and fin count were used in this study rather than giving a range of values in a smaller interval. This was done to restrict the design space to a certain number of design points and reduce simulation time. As expected, the pressure drop increases and thermal resistance reduces with the fin thickness of the heat sink. It was, however, observed that the variation in thermal resistance becomes less than 1% after a fin thickness of 0.6 mm. The same variation between 0.25 and 0.6 mm fin thickness is around 17%. This means a final optimized design of the heat sink will require a trade-off between the pressure drop and thermal resistance values.

To observe the coupled effect of more than one design variable on the objective function, we use 3D response surface plots. Figure 9 shows the response plots for the variation of two objective functions with heat-sink fin thickness and fin height. It is clear from the response plots that both the objective functions have an inverse relationship with each other. This would mean that to achieve better thermal performance, a pumping-power penalty would be incurred at some point. The CoP values of 100% were obtained for all the response surfaces of both the objective functions which depict the quality of the model approximation of the input design space. The optimization solver uses the design space created by the sensitivity analysis and plots the




FIG. 8: Relation between thermal resistance and different input variables in 2D regression plot for the aluminum heat sink in forced convection

solution of the objective functions as outputs on a Pareto front as shown in Fig. 10. Each point on this Pareto front represents a solution for the objective functions based on a design space point. The best design points from the Pareto front are those that lie along the red boundary. It can be seen that a vast majority of the design points were clustered in a region with a pressure drop of less than 20 Pa and a thermal resistance value between 0.2 and 0.18 W/mK. After analyzing all of the best design points (DPs), five design points have both a lower pressure drop and thermal resistance as shown in Table 3. Out of the selected points, DP 163 showed the most optimum value for pressure drop, source temperature, and thermal resistance. When compared to the baseline heat-sink design, the optimized heat sink offers a 15.3% less pressure drop and approximately 15% lower thermal resistance for the same fin height.

In the data center industry, server heights are typically denoted in terms of form factor. A form factor of 1U represents a server height of 44.5 mm. As discussed earlier, one of the main advantages of immersion cooling lies in the fact that it allows a higher heat-transfer rate as compared to air-cooling technologies. Therefore, it is also possible to reduce the form factor




FIG. 9: Comparison of 3D response surfaces of the objective functions with varying fin thickness and heat-sink height




FIG. 10: Pareto front showing the entire design space of 216 points and the red boundary depicting the best design points where the objective functions have a minimum value

or height of the heat sink itself to dissipate the same amount of heat flux from the processors. This allows the data center to pack more servers in the same space by utilizing efficient cooling methods like immersion cooling. The baseline heat sink used in this study is designed for a 2U

TABLE 3: Summary of best design points showing the values of corresponding objective functions and source temperature

Design point	HS height	Fin count	Fin thickness	R_{th}	ΔP	T_{source}
Baseline	41	35	0.23	0.22	11.3	57.4
35	41	25	0.5	0.19	9.6	54.8
131	41	25	0.68	0.18	10.2	53.3
136	41	27	0.41	0.20	9.2	55.9
143	41	27	0.68	0.18	11	52.9
163	41	25	0.59	0.19	9.9	53.9

air-cooled server. An investigation of reducing the heat-sink height was also carried out from the optimization results. To do a fair comparison with the baseline design, the selection criterion was which heat sink would give a lower temperature than the baseline design for a reduced heat-sink height. The optimization results of the heat sink with reduced height show that the best DPs for a heat-sink height equivalent to a 1U server, the optimized heat sink offers lower thermal resistance and source temperature than the baseline design in a lower form factor for the same CPU power. This means that while the overall heat-transfer area is being reduced, an optimized heat sink for immersion cooling offers better thermal performance as seen in Table 4, although the pressure drop or the pumping power will increase significantly for a smaller form-factor server. The trade-off to be weighed here is the improvement in thermal performance and the ability to deploy more servers in the same volume of space in a data center. Further analysis of the reduced form-factor heat sink also shows that a significant reduction in heat-sink weight and production cost per unit can also be reduced.

As discussed in Section 3, after determining the optimized heat-sink geometry for the baseline design, an optimization study was also done for natural convection immersion cooling. Figure 11 shows the total effects plot for the aluminum heat sink in natural convection. It is seen that unlike the case of forced convection, which is still very low Reynolds number flow in immersion cooling, the impact of fin thickness on thermal resistance reduces from 82 to 59%. Another difference observed when comparing the case of natural convection to forced convection is that the impact of heat-sink height becomes at least 5 times more significant. At the same time, the effect of heat-sink fin count becomes even less important in terms of optimization for natural convection. For pressure drop, it is observed that the impact number of fins increases almost threefold and the impact of overall heat-sink height reduces by almost 50%.

TABLE 4: Overview of the results of design points with heat sink height for a 1U server

Design point	HS height	Fin count	Fin thickness	R_{th}	ΔP	T_{source}
Baseline	41	35	0.23	0.22	11.3	57.4
120	26	35	0.68	0.17	28.1	51.9
99	26	25	0.32	0.20	17.5	56.3
41	26	27	0.23	0.21	18.2	57.2
171	26	31	0.68	0.17	25.7	51.7
24	26	33	0.41	0.18	24.1	56.5




FIG. 11: Total effects plot for aluminum heat sink under natural convection

The distribution of the solution for each DP on the Pareto front is shown in Fig. 12. This is also different when compared to Fig. 10, where a greater number of solutions satisfy the objective function constraints and are more uniformly distributed. As compared to the results of pressure drop and thermal resistance of the baseline design in natural convection, a maximum reduction of 32% was obtained in pressure drop and a 10% reduction in thermal resistance. Additionally, the optimized heat sink was able to reduce the source temperature by 3°C for the same fin height as the baseline design. An interesting observation made from the response graphs is that the thermal resistance value reduces to a certain fin count, and then starts creeping up as seen in Fig. 13. This could be because a larger number of fins tend to choke the natural convection through the fins. This effect is also observed in thermal resistance response to fin thickness that plateaus out after 0.5 mm fin thickness.

4.2 Optimization for Copper Heat Sink

Aluminum is typically used for heat-sink applications as it offers lower weight and cheaper manufacturing options. Copper heat sinks offer better thermal performance, especially when the application is for high-power density. Therefore, following the same methodology as for aluminum heat sinks, a design optimization was also done for the same heat sink with copper as the material. Both the design variables and the objective functions were kept the same as in the case of aluminum heat sinks. A baseline CFD simulation was first carried out to ascertain the values of the objective functions under both forced and natural convection cooling modes. Figure 14 shows the total effects matrix for the copper heat sink DoE. It can be seen from the matrix that




FIG. 12: Pareto front showing the distribution of the solutions of the objective functions and the chosen DPs




FIG. 13: Response graph for variation of thermal resistance with the design points for heat-sink fin count

fin thickness and fin height have a dominant impact on thermal resistance and pressure drop, respectively. The Pareto front distribution looks very similar to the forced convection case in the aluminum heat sink. A comparison of the objective function values of the best DPs from the Pareto front is shown in Table 5.

Also, as done in the case of the aluminum heat sink, an analysis was done to determine the performance of the optimized heat-sink design for the same source power but a smaller 1U form




FIG. 14: Total effects matrix for copper heat sink with forced convection cooling

TABLE 5: Summary of the results of the optimized values of the objective function for forced convection cooling using copper heat sinks

Design point	HS height	Fin count	Fin thickness	R_{th}	ΔP	T_{source}
Baseline	41	35	0.23	0.15	11.3	49.6
27	41	27	0.68	0.13	11	47.3
28	41	25	0.5	0.13	9.5	48.2
158	41	25	0.41	0.13	10.2	47.4
199	41	27	0.68	0.14	9.2	48.8
207	41	25	0.59	0.13	9.2	47.7

factor. As opposed to the case of the aluminum heat sink where the pressure drop increases significantly, an optimized heat-sink design for copper showed negligible variation in pressure drop is observed. At the same time, the optimized 1U heat sink displayed the same thermal performance as the baseline heat sink as seen in Table 6.

For the case of natural convection-cooled copper heat sinks, the main difference in the total effects plot is in pressure drop in both the number of fins and fin height. Also, as opposed to forced convection, thermal resistance for natural convection copper heat sink is highly dependent


TABLE 6: Results of the optimized values of the objective functions for 1U copper heat sink

Design point	HS height (mm)	Fin count	Fin thickness (mm)	R_{th} (W/mK)	ΔP (Pa)	T_{source} (°C)
Baseline	41	35	0.23	0.15	11.3	49.6
27	26	25	0.59	0.15	11.3	49.6

on fin height as shown in Fig. 15. This is also different from the aluminum heat sink in natural convection where the thermal resistance is dominated by a variation in fin thickness. The Pareto front for the solution of the objective functions shows that the best DP for the heat-sink design reduces the thermal resistance by approximately 6% and reduces the pressure drop by 27% as compared to the baseline design of copper heat sink in natural convection. However, the copper heat sink will add extra weight and in such a case, a trade-off needs to be considered in terms of cost and mechanical constraints.

5. CONCLUSION

As the power densities for high-performance processors keep rising, the demands for the deployment of efficient cooling technologies are also increasing. Single-phase immersion cooling helps to address many shortcomings of traditional air-cooling and competitive liquid-cooling

**FIG. 15:** Total effects plot for copper heat sink under natural convection

technologies. As compared to air-cooling, immersion-cooling offers higher thermal mass, simplicity in cooling infrastructure, alleviation of airborne contamination issues, and is particularly suitable for edge data center deployments. An in-depth study of different multi-objective and multi-design variable optimization schemes for heat sinks in immersion-cooled servers was investigated. Heat-sink geometric parameters like heat-sink fin height, fin thickness, and fin count were varied. Minimization of thermal resistance and pressure drop for constant pumping power was chosen as the objective function for the optimization study. As compared to the baseline heat-sink design, the optimized heat sink was able to reduce the thermal resistance by 15% and a further 15.3% reduction in pumping power. A comparison of the optimization results from the aluminum heat sink between forced and natural convection shows that the impact of fin thickness reduces significantly and the influence of heat-sink height increases almost five times. This change of influence is even more pronounced for copper heat sink where the influence of heat-sink height is negligible in forced convection but has the most dominant influence in natural convection.

The results of this study on the influence of the heat-sink design variables on heat transfer and pumping power can be directly used for the optimization of typical parallel plate heat sinks for single-phase immersion cooling. There are more heat-sink fin designs and configurations that can also be studied for optimization. As an example, pin-fin heat sinks are typically used in natural convection flows but offer a lower overall heat-transfer surface area. An optimized pin-fin geometry can be obtained that not only offers lower pressure drop but also enhances their heat-transfer performance. A further extension of this study can be to carry out a multi-disciplinary optimization where the optimization parameters can be thermal and flow physics, a cost model of the heat sink, and material optimization based on the thermal profile at the heat-sink base.

REFERENCES

- Bansode, P.V., Shah, J.M., Gupta, G., Agonafer, D., Patel, H., Roe, D., and Tufty, R., Measurement of the Thermal Performance of a Single-Phase Immersion Cooled Server at Elevated Temperatures for Prolonged Time, in *Proc. of the ASME 2018 Intl. Technical Conf. and Exhibition on Packaging and Integration of Electronic and Photonic Microsystems*, San Francisco, CA, August 27–30, 2018.
- Bansode, P.V., Shah, J.M., Gupta, G., Agonafer, D., Patel, H., Roe, D., and Tufty, R., Measurement of the Thermal Performance of a Custom-Build Single-Phase Immersion Cooled Server at Various High and Low Temperatures for Prolonged Time, *ASME J. Electron. Packag.*, vol. **142**, no. 1, Article ID 011010, 2020.
- Brink, R., et al., Design Guidelines for Immersion-Cooled IT Equipment, *Open Compute Project Advanced Cooling Solutions Immersion Workstream*, accessed March 2022, from <https://docs.google.com/document/d/1VN0yd5bySgiIxxWTpWUcOejonp6cz5D/edit#>, 2022.
- Chen, C.T. and Chen, H.I., Multi-Objective Optimization Design of Plate-Fin Heat Sinks Using a Direction-Based Genetic Algorithm, *J. Taiwan Inst. Chem. Eng.*, vol. **44**, no. 2, pp. 257–265, 2013.
- Chen, C.T., Wu, C.K., and Hwang, C., Optimal Design and Control of CPU Heat Sink Processes, *IEEE Trans. Compon. Packag. Technol.*, vol. **31**, no. 1, pp. 184–195, 2008.
- Chiang, K.T., Chang, F.P., and Tsai, T.C., Optimum Design Parameters of Pin-Fin Heat Sink Using the Grey-Fuzzy Logic Based on the Orthogonal Arrays, *Int. Commun. Heat Mass Transf.*, vol. **33**, no. 6, pp. 744–752, 2006.
- Chiang, K.T. and Chang, F.P., Application of Response Surface Methodology in the Parametric Optimization of a Pin-Fin Type Heat Sink, *Int. Commun. Heat Mass Transf.*, vol. **33**, no. 7, pp. 836–845, 2006.

- Devi, S.P., Manivannan, S., and Rao, K.S., Comparison of Nongradient Methods with Hybrid Taguchi-Based Epsilon Constraint Method for Multiobjective Optimization of Cylindrical Fin Heat Sink, *Int. J. Adv. Manufact. Technol.*, vol. **63**, nos. 9-12, pp. 1081–1094, 2012.
- Eiland, R., Fernandes, J., Vallejo, M., Agonafer, D., and Mulay, V., Flow Rate and Inlet Temperature Considerations for Direct Immersion of a Single Server in Mineral Oil, in *14th Intersociety Conf. on Thermal and Thermomechanical Phenomena in Electronic Systems (ITherm)*, Lake Buena Vista, FL, pp. 706–714, 2014.
- Engineered Fluids, *EC-100 Dielectric Coolant*, accessed September 3, 2021, from <https://engineeredfluids.store/products/ec-100>, 2021.
- Gandhi, D., Chowdhury, U., Chauhan, T., Bansode, P., Saini, S., Shah, J.M., and Agonafer, D., Computational Analysis for Thermal Optimization of Server for Single Phase Immersion Cooling, in *Proc. of the ASME 2019 Intl. Technical Conf. and Exhibition on Packaging and Integration of Electronic and Photonic Microsystems*, Anaheim, CA, October 7–9, 2019.
- Greenberg, S., Mills, E., Tschudi, B., Rumsey, P., and Myatt, B., Best Practices for Data Centers: Lessons Learned from Benchmarking 22 Data Centers, in *Proc. of the ACEEE Summer Study on Energy Efficiency in Buildings*, Asilomar, CA, pp. 76–87, August 3, 2006.
- 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, *IEEE Trans. Compon. Packag. Manufact. Technol.*, vol. **11**, no. 10, pp. 1565–1582, 2021.
- Internet Live Stats, accessed February 1, 2022 from <https://www.internetlivestats.com/>, 2022.
- 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.
- Kim, D.K., Thermal Optimization of Plate-Fin Heat Sinks with Fins of Variable Thickness under Natural Convection, *Int. J. Heat Mass Transf.*, vol. **55**, no. 4, pp. 752–761, 2012.
- Masanet, E., Shehabi, A., Lei, N., Smith, S., and Koomey, J., Recalibrating Global Data Center Energy-Use Estimates, *Science*, vol. **367**, no. 6481, pp. 984–986, 2020.
- McWilliams, T.D., Evaluating Heat Sink Performance in an Immersion-Cooled Server System, MS, The University of Texas at Arlington, USA, 2014.
- Mimery, D.R., Multidisciplinary Design Optimization of Part Geometry in CAD, MS, Massachusetts Institute of Technology, USA, 2020.
- Most, T. and Will, J., Sensitivity Analysis Using the Metamodel of Optimal Prognosis, in *8th Optimization and Stochastic Days*, Weimar, Germany, pp. 24–40, November 24–25, 2011.
- Niazmand, A., Murthy, P., Saini, S., Shahi, P., Bansode, P., and Agonafer, D., Numerical Analysis of Oil Immersion Cooling of a Server Using Mineral Oil and Al_2O_3 Nanofluid, in *Proc. of the ASME 2020 Intl. Technical Conf. and Exhibition on Packaging and Integration of Electronic and Photonic Microsystems*, Virtual, October 27–29, 2020a.
- Niazmand, A., Chauhan, T., Saini, S., Shahi, P., Bansode, P.V., and Agonafer, D., CFD Simulation of Two-Phase Immersion Cooling Using FC-72 Dielectric Fluid, in *Proc. of the ASME 2020 Intl. Technical Conf. and Exhibition on Packaging and Integration of Electronic and Photonic Microsystems*, Virtual, October 27–29, 2020b.
- Saini, S., Shah, J. M., Shahi, P., Bansode, P., Agonafer, D., Singh, P., Schmidt, R., and Kaler, M., Effects of Gaseous and Particulate Contaminants on Information Technology Equipment Reliability—A Review, *ASME J. Electron. Packag.*, vol. **144**, no. 3, Article ID 030801, 2022.
- Shah, J.M., Eiland, R., Siddarth, A., and Agonafer, D., Effects of Mineral Oil Immersion Cooling on IT Equipment Reliability and Reliability Enhancements to Data Center Operations, in *2016 15th IEEE Intersociety Conf. on Thermal and Thermomechanical Phenomena in Electronic Systems (ITherm)*, Las

- Vegas, NV, pp. 316–325, May 31–June 3, 2016.
- Shah, J.M., Eiland, R., Rajmane, P., Siddarth, A., Agonafer, D., and Mulay, V., Reliability Considerations for Oil Immersion-Cooled Data Centers, *ASME J. Electron. Packag.*, vol. **141**, no. 2, Article ID 021007, 2019.
- Shah, J.M., Bhatt, C., Rachamreddy, P., Dandamudi, R., Saini, S., and Agonafer, D., Computational Form Factor Study of a 3rd Generation Open Compute Server for Single-Phase Immersion Cooling, in *Proc. of the ASME 2019 Intl. Technical Conf. and Exhibition on Packaging and Integration of Electronic and Photonic Microsystems*, Anaheim, CA, October 7–9, 2019.
- Shah, J.M., Dandamudi, R., Bhatt, C., Rachamreddy, P., Bansode, P., and Agonafer, D., CFD Analysis of Thermal Shadowing and Optimization of Heatsinks in 3rd Generation Open Compute Server for Single-Phase Immersion Cooling, in *Proc. of the ASME 2019 Intl. Technical Conf. and Exhibition on Packaging and Integration of Electronic and Photonic Microsystems*, Anaheim, CA, October 7–9, 2019.
- Shah, J.M., Padmanaban, K., Singh, H., Duraisamy Asokan, S., Saini, S., and Agonafer, D., Evaluating the Reliability of Passive Server Components for Single-Phase Immersion Cooling, *ASME J. Electron. Packag.*, vol. **144**, no. 2, Article ID 021109, 2022.
- Shahi, P., Agarwal, S., Saini, S., Niazmand, A., Bansode, P., and Agonafer, D., CFD Analysis on Liquid Cooled Cold Plate Using Copper Nanoparticles, in *Proc. of the ASME 2020 Intl. Technical Conf. and Exhibition on Packaging and Integration of Electronic and Photonic Microsystems*, Virtual, October 27–29, 2020.
- Shahi, P., Saini, S., Bansode, P., and Agonafer, D., A Comparative Study of Energy Savings in a Liquid-Cooled Server by Dynamic Control of Coolant Flow Rate at Server Level, *IEEE Trans. Compon. Packag. Manuf. Technol.*, vol. **11**, no. 4, pp. 616–624, 2021.
- 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**, no. 4, Article ID 041008, 2022.
- Shehabi, A., Smith, S.J., Masanet, E., and Koomey, J., Data Center Growth in the United States: Decoupling the Demand for Services from Electricity Use, *Environ. Res. Lett.*, vol. **13**, no. 12, p. 124030, 2018.
- Shinde, P.A., Bansode, P.V., Saini, S., Kasukurthy, R., Chauhan, T., Shah, J.M., and Agonafer, D., Experimental Analysis for Optimization of Thermal Performance of a Server in Single-Phase Immersion Cooling, in *Proc. of the ASME 2019 Intl. Technical Conf. and Exhibition on Packaging and Integration of Electronic and Photonic Microsystems*, Anaheim, CA, October 7–9, 2019.
- Siddik, M.A.B., Shehabi, A., and Marston, L., The Environmental Footprint of Data Centers in the United States, *Environ. Res. Lett.*, vol. **16**, no. 6, Article ID 064017, 2021.
- Subasi, A., Sahin, B., and Kaymaz, I., Multi-Objective Optimization of a Honeycomb Heat Sink Using Response Surface Method, *Int. J. Heat Mass Transf.*, vol. **101**, pp. 295–302, 2016.
- Will, J. and Most, T., Metamodell of Optimized Prognosis (MoP)—An Automatic Approach for User Friendly Parameter Optimization, in *6th Optimization and Stochastic Days*, Weimar, Germany, October 15–16, 2009.
- Yazicioglu, B. and Yüncü, H., Optimum Fin Spacing of Rectangular Fins on a Vertical Base in Free Convection Heat Transfer, *Heat Mass Transf.*, vol. **44**, no. 1, pp. 11–21, 2007.