

# Hydraulic and Thermal Characteristics of a Double-Sided Cold Plate – Part 2

## Tesla Tear Down, CFD Validation, and Machine Learning to Determine the Performance Limit

**Azita Soleymani**  
HeatSync

**John Wilson**  
Siemens Digital Industries Software

**William Maltz**  
Electronic Cooling Solutions, Inc.

This study investigates the hydraulic and thermal characteristics of the TESLA AUTOPILOT HW2.5 MODEL 3 Y, which features a double-sided cold plate with PCBs attached on both sides. We present a step-by-step teardown process of the unit, measuring internal dimensions, fin locations, inlet/exhaust ports, pedestals, and cold-plate channels. These measurements enabled the creation of a 3D numerical

model for the Tesla AUTOPILOT HW2.5 unit, which was used to analyze the hydraulic and thermal performance of the cold plate.

We conducted experiments to measure hydraulic head loss through the cold plate and evaluate thermal performance at various coolant flow rates and power loads. These results were used to validate the CFD model. Thermal loads and the fluid flow rate



**Azita Soleymani**

Dr. Azita Soleymani is the CTO of HeatSync, empowering excellence through thermal innovation in electronics and battery systems. With nearly 150 years of combined expertise, HeatSync (<https://heat-sync.com/>) tackles some of the most demanding thermal challenges across industries, including consumer electronics, automotive, energy storage, telecommunications, aerospace, medical devices, data centers, AI systems, and more. Formerly a member of the technical staff at Meta, Dr. Soleymani has co-authored over 40 technical papers and is a recognized leader in thermal design, simulation, testing, and optimization.



**John Wilson**

*Electronics Thermal Application Specialist at Siemens, Digital Industries Software.* John Wilson joined Mentor Graphics Corporation, Mechanical Analysis Division (formerly Flomerics Ltd), that is now part of Siemens Digital Industries Software, in 1999 and has over 25 years of thermal design experience relating to simulation and testing. John currently works with product management teams to provide electronics thermal design solutions to leading electronics industry clients globally.



**William Maltz**

William Maltz has 38 years of experience. The ECS team, he leads, provides design advice to clients. ECS uses numerical analysis and experimental work to identify design solutions. He chaired technical sessions at SemiTherm, IMAPS, InterPACK, ITherm and co-authored technical papers and magazine articles on CFD, liquid cooling, natural convection, optimization techniques, and boundary condition independent models. He has reviewed books, publications, and presentations. William mentored 20+ engineers, who utilized their ECS experience to launch careers elsewhere. He is an ASME Fellow and active in the local ASME Section.

in the CFD model mimicked the experimental test cases that were evaluated. The model results for the hydraulic pressure loss and temperatures at specific locations were found to correlate well with the experimental data that was collected. The validated CFD model was then used to examine the thermal characteristics of the cold plate across a range of operating conditions. The model was also used to identify opportunities to improve the design of the cold plate.

Additionally, an Artificial Neural Network (ANN) model of the cold plate assembly was developed to analyze system performance across a *wide range* of operating parameters. It was demonstrated that the developed neural network model could be used to determine the performance limit of the cold plate.

**Keywords:** Cold plate, Autonomous driving, CFD, Tear down, Thermal performance, Hydraulic performance, Artificial Neural Network, Machine Learning.

### Introduction

Part 1 of this paper [1] detailed the construction and simulation results of the CFD model of the cold plate assembly, providing insights into its performance through velocity and temperature profiles. Part 2 outlines the teardown process and experimental setup used to validate the simulation model.

Artificial intelligence (AI) and machine learning (ML) have gained prominence in both scientific and public discourse over the past decade. ML algorithms utilize data subsets to generate predictive rules for system outcomes based on input variables. The ANN, a subset of ML, learns from example data to form probability-weighted associations between inputs and results that are stored within its structure.

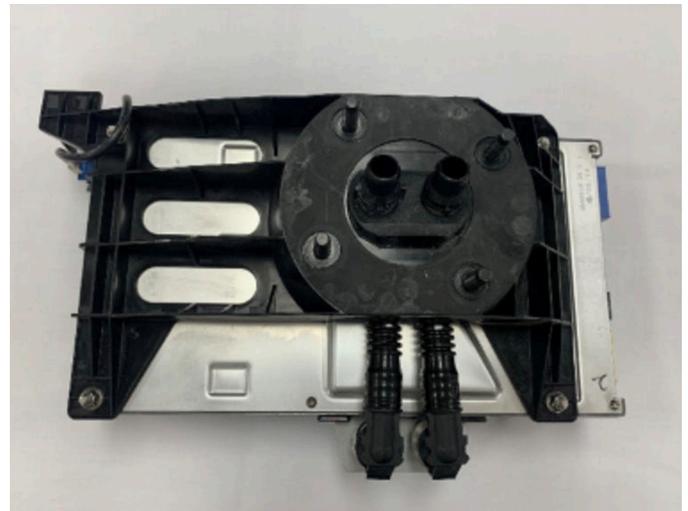
In this paper, an ANN of the cold plate assembly is presented to study the performance of the system using multiple combinations of the operating parameters.

### Tesla Autopilot Assembly Tear Down

The autopilot unit, shown in *Figure 1*, is enclosed in an aluminum housing and with a cold plate transition assembly attached to the front side of the housing.

First, the autopilot cold plate transition assembly, which includes the coolant hoses shown in *Figure 2*, was detached from the housing of the autopilot assembly. The hoses are essential components of the cooling loop, allowing coolant to flow into and out of the cold plate. Additionally, *Figure 2* illustrates the attachment points for the hoses on the transition assembly, providing a clear view of the pathways for coolant circulation within the system.

Then, the top and bottom sections of the autopilot aluminum housing were removed. As shown in *Figure 3*, the cold plate serves as the central structure within the assembly, with printed circuit boards (PCBs) securely mounted on both sides.



Front



Back

Figure 1: The autopilot assembly before teardown

As shown in *Figure 4*, these PCBs house critical components, including four high-power GPUs or processors, which are strategically positioned for optimized thermal management. Additionally, various support electronics and connectors are visible on the PCBs, designed to facilitate signal processing and power management for the autopilot system.

The cold plate, after the thermal putty was removed, is shown in *Figure 5*. Finned channels are integrated into the cold plate design to improve the heat transfer. There are four lapped, or polished, pedestals with the four main chips attached to them. These pedestals are labeled A, B, C, and D to provide reference points for the discussion in this document.

The cold plate was dismantled so that the dimensions of the internal geometry could be measured. The measurements were used to generate a CAD model of the unit.

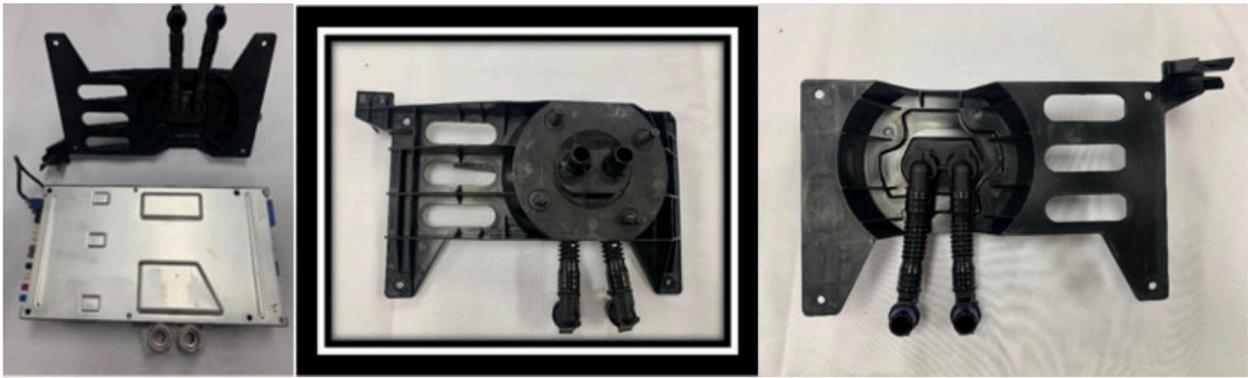


Figure 2: The hoses are attached to the autopilot cold plate transition assembly

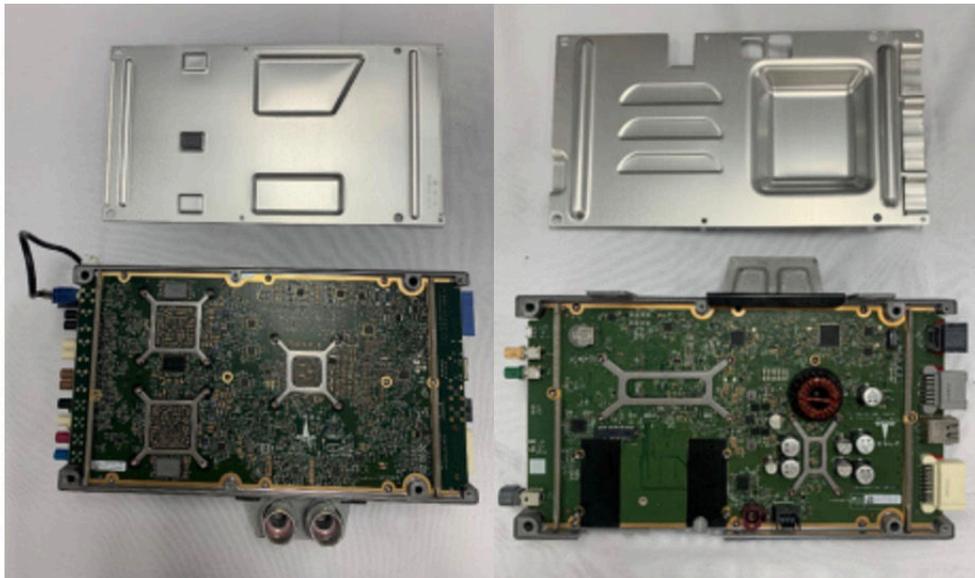


Figure 3: The autopilot consists of two PCBs attached to both sides of a cold plate

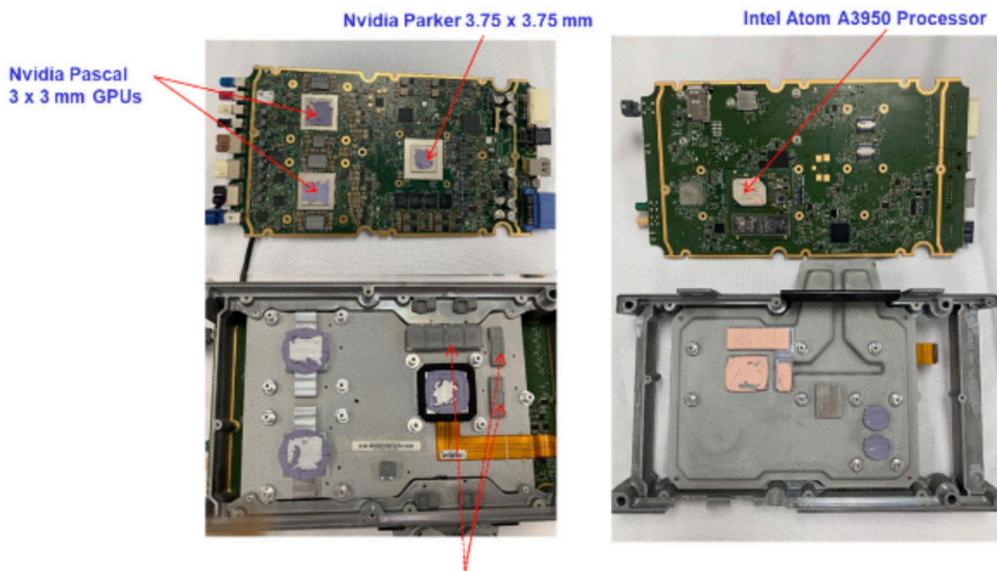


Figure 4: Removal of the PCBs from the Cold Plate

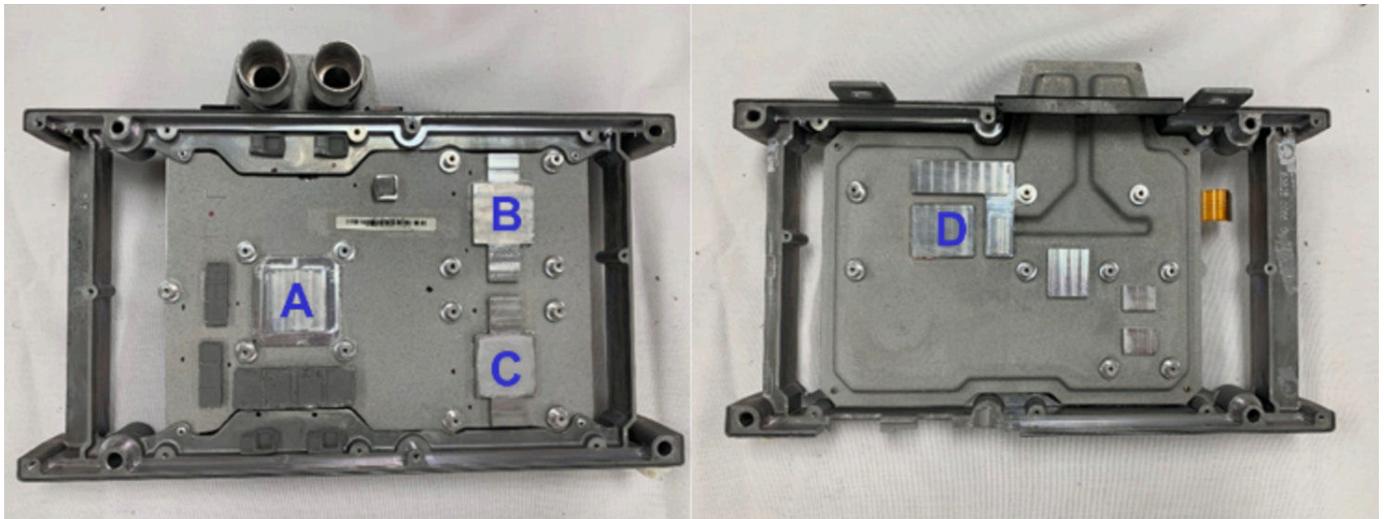


Figure 5: Front and Back Views of the Die Casted Cold Plate

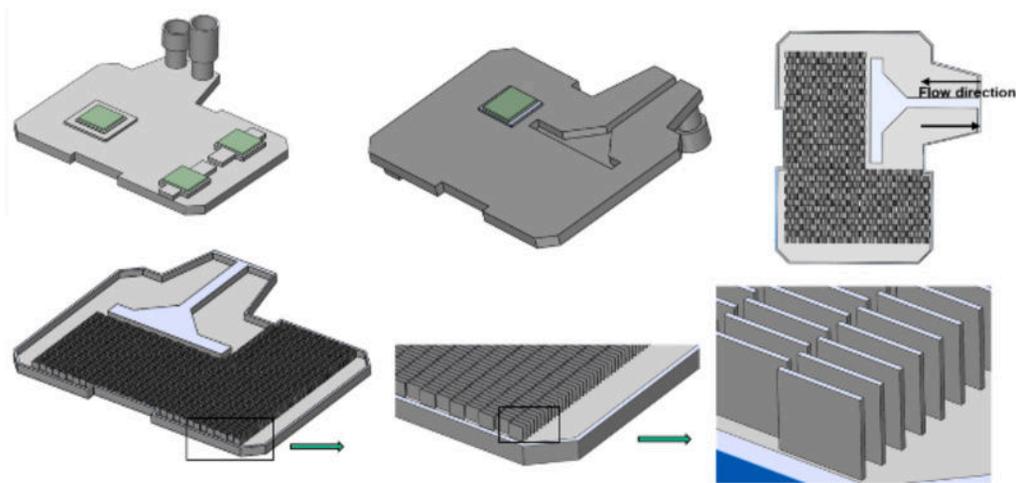


Figure 6: The CAD Model Representing the Double-sided Cold Plate

### Experimental Setup

A set of experiments were conducted to measure the hydraulic head loss through the cold plate and assess its thermal performance under varying coolant flow rates and power loads. To facilitate these tests, a second autopilot unit was disassembled, and thermocouples and heaters representing the four main chips were attached to the cold plate. The coolant used in these experiments was a 50/50 glycol-water mixture, consistent with the specifications provided by the original design. This coolant choice aligns with the intended operating conditions, ensuring that test results accurately reflect the cold plate's performance in real-world applications.

Temperatures of the pedestal below each heater module pressure head loss through the cold plate, and the liquid temperatures were all measured. Uniform power was applied to the four heaters.

The heaters mounted to the cold plate were driven by a variable

power supply so that a specified power dissipation could be applied to the cold plate. The four heaters were placed in the locations where the GPUs and processors contact the cold plate. *Figure 7* shows locations where thermocouples and heaters were attached to the front side of the cold plate.

Arctic Silver compound was placed on the interfaces between the heaters and the cold plate pedestals. A 2 mm x 1 mm groove was milled from the side to the center of each pedestal. The thermocouple beads were then secured in the groove with Arctic Alumina epoxy. After curing, the grooves were filled with additional Arctic Alumina epoxy, and a wood insulator was placed over each heater to improve thermal control. The entire assembly was then thermally insulated; however, while insulating the heaters themselves enhanced measurement accuracy, insulating the full assembly had a negligible effect on accuracy within the range studied.

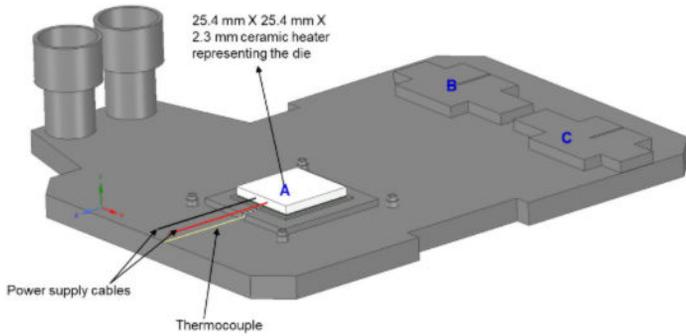


Figure 7: Instrumenting the cold plate; to provide better visualization, just one heater assembly is shown

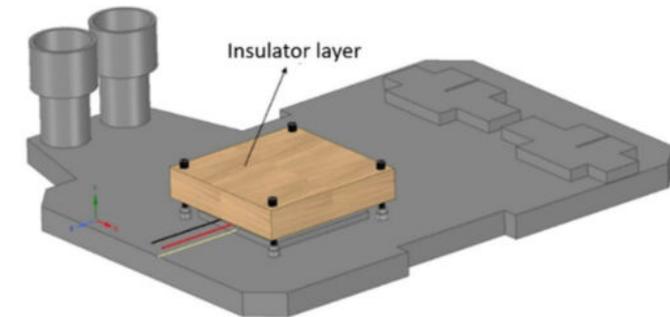
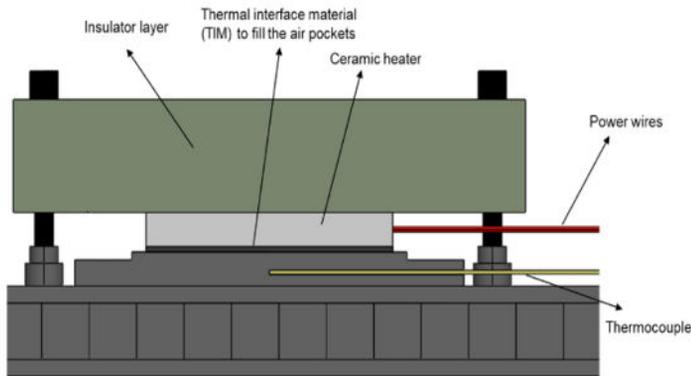


Figure 8: Heater module assembly

The schematic of the cooling loop is illustrated in Figure 9. All the sensors were calibrated prior to use.

Two thermocouples were placed inside the liquid loop to measure the coolant temperature upstream and downstream of the cold plate. To ensure accurate coolant outflow temperature readings, a custom-designed, 3D-printed static mixer was placed upstream of the thermocouples. Energy balance was achieved by calculating the heat generated by the heaters and comparing it to the energy extracted by the coolant in steady-state mode, resulting in a difference of less than 1% of the total heater energy. The orientation of the cold plate during testing did not impact hydraulic or thermal performance. Figure 10 provides a snapshot of the experimental setup.

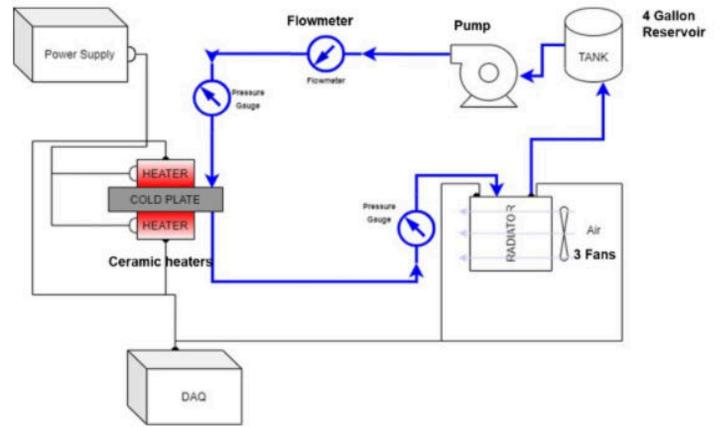


Figure 9: Simplified cooling loop schematic

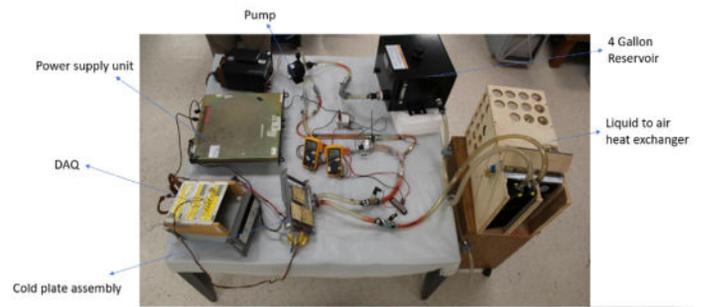


Figure 10: Overall experimental set-up

### Validation of the CFD model

To validate the numerical simulations, test results for the hydraulic pressure drop across the cold plate are compared to simulations in Figure 11. The CFD model for hydraulic pressure loss correlated with the experimental data over a wide range of operating conditions. The minimum and maximum pressure loss of the unit were measured to be 0.4 kPa and 10 kPa at 0.4 and 6.4 LPM, respectively.

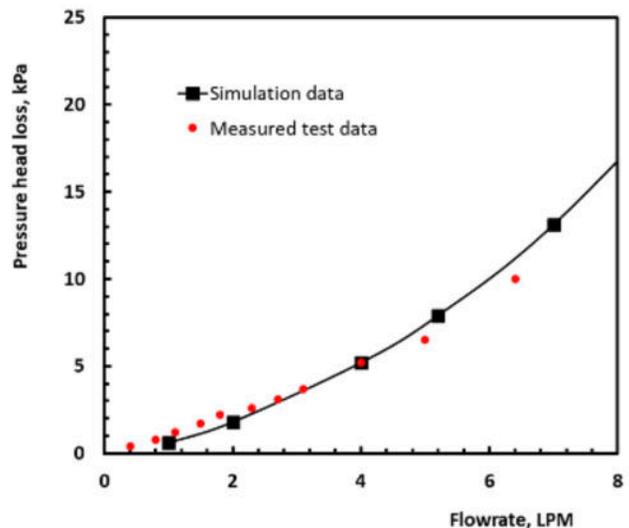


Figure 11: The pressure drop across the cold plate

	Pedestal temperature, °C					
	Case1 Coolant flow rate: 5.2 LPM			Case2 Coolant flow rate: 4 LPM		
	Simulation	Test data	$T_{CFD}-T_{experiment}$	Simulation	Test data	$T_{CFD}-T_{experiment}$
Chip A	45.3	46	-0.7	46	48	-2
Chip B	47.9	48	-0.1	48.9	50	-1.1
Chip C	47.3	48	-0.7	48.6	50	-1.4
Chip D	49.7	51	-1.3	51.1	54	-2.9

Table 1: Simulations vs. experiment results for coolant with an inlet temperature of 30°C, heat dissipation rate of 125.5 W per chip

The pedestal temperature test results of the four chips are compared to the simulation results in *Table 1*. This shows that the modeling results for the temperature map were consistent with the experimental data with a maximum temperature difference of less than 3°C.

The higher the flow rate, the lower the temperature of the chips. However, a more powerful pump is required to provide the coolant at higher flow rates. There is a tradeoff between the performance of the cold plate, the pump size, and the energy required to push the coolant through the cold plate. The liquid flow rate should be optimized to maintain a relatively low temperature rise from inlet to exhaust while also minimizing the impact on system pressure drop and pumping requirements.

#### Artificial Neural Network Model

The Artificial Neural network was organized in layers. The layers were made up of many interconnected *nodes* that contained *activation functions*. The number of the nodes in an input layer was set to 6 to include coolant inflow temperature, flow rate and thermal loads of four IC components. The input layer contained the values of the explanatory attributes for each observation. The hidden layers apply given transformations to the input values inside the network. For the cold plate assembly, two hidden layers were defined. The first and second layers included 12 and 6 hidden nodes, respectively. The response variables in the output layer included hydraulic pressure loss and the junction temperature for the four IC components. In this analysis, the temperatures of the pedestals beneath each heater module were used to represent the junction temperatures. This approach allowed for consistent temperature measurement relevant to the thermal management of the ICs.

The learning set was obtained by performing a set of paramet-

ric CFD simulations in which the coolant viscosity varied with temperature. Design of Experiments (DOE) was conducted, applying Latin Hypercube Sampling method in the design space. The coolant flow rate and temperature varied between 2 LPM to 7 LPM and 6°C to 44°C, respectively. The thermal load of each IC was set between 0 to 200 W. In total, the DOE included 140 simulation points, with 120 of them randomly selected as the learning set, to avoid overfitting, and the remaining 20 points were used to validate the generated ANN model. The predicted results for the hydraulic pressure loss across the cold plate and the four junction temperatures were compared to those from sample set. The correlation coefficient ( $R^2$ ) for both the pressure loss and the junction temperatures was calculated to be over 0.98. The ANN results were found to correlate well with the validated CFD results. The ANN model can be used to calculate the hydraulic and thermal performance of the cold plate at any operating conditions within seconds without compromising accuracy.

To show case the capabilities of the generated ANN model, it was applied to estimate the thermal capacity of the cold plate, assuming an upper junction temperature limit of 90°C and a coolant inlet temperature of 40°C. As shown in *Figure 12*, the current cold plate can remove a maximum of 155 W and 179 W from Chip A at a coolant flow rate of 2 LPM and 7 LPM, respectively.

The power dissipation rate of Chip A increases the junction temperature of chip B, C, and D by a maximum of 7°C, 4.6°C and 6.6°C at 2 LPM and 1.6°C, 0.4°C and 1.6°C at 7 LPM, respectively. Assuming the maximum heat dissipation rate for Chip A, the cold plate can remove a maximum of 125 W and 156 W from Chip C at 2 LPM and 7 LPM, respectively.

The cold plate has the capacity to remove a maximum of 110 W and 156 W from Chip B at 2 LPM and 7 LPM, respectively, as-

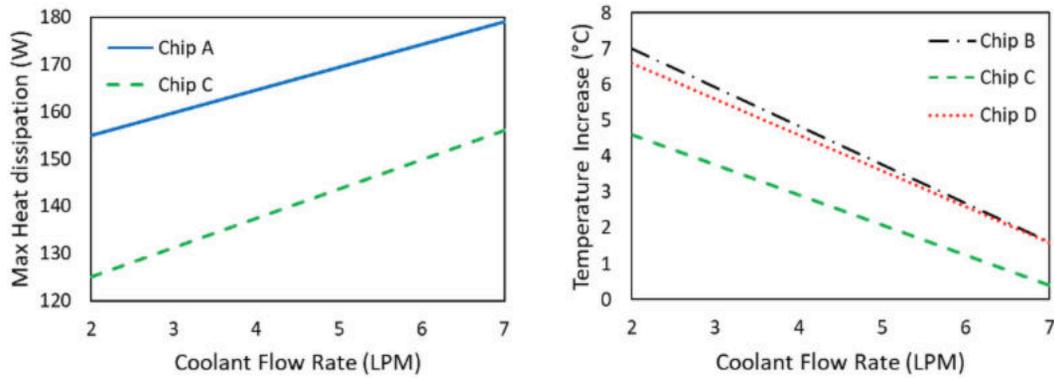


Figure 12: Maximum heat dissipation for Chips A and C at various coolant flow rates, along with the junction temperature increases for Chips B, C, and D due to Chip A's dissipation, based on an upper junction temperature limit of 90°C and a coolant inlet temperature of 40°C.

suming the maximum heat load from chip A and C. The maximum heat loads of Chip D were found to be similar to Chip B, 113 W and 155 W at 2 LPM and 7 LPM, respectively. The total thermal capacities of the cold plate were 503 W and 646 W at 2 LPM and 7 LPM.

### Conclusion

It was shown that the developed artificial neural network model can represent the thermal and hydraulic characteristics of the cold plate. The results for pressure head loss and junction tem-

peratures align closely with validated CFD results across a range of conditions. The model facilitates temperature estimation for critical components and evaluates cold plate performance under various coolant inflow conditions and power load distributions. The generated ANN model was used to identify performance limits of the cold plate. The developed model can accelerate cold plate design and development by expediting an evaluation of design feasibility and to conduct in-depth root cause analyses for various inputs and operating conditions.

### References

- [1] Azita Soleymani et al., "Hydraulic and Thermal Characteristic of a Double-Sided Cold Plate. Part 1: CFD Analysis", *Electronics Cooling Magazine*, Spring 2023.