NEXALUS
COOLING THE CLOUD

Closed Loop Liquid Cooling of High-Powered CPUs: A Case Study on Cooling Performance and Energy Optimization

F.M. Naduvilakath-Mohammed[a,b,c], R. Jenkins[b], G. Byrne[a], A.J. Robinson[a,b,c]

a Department of Mechanical and Manufacturing Engineering, Parsons Building, Trinity College Dublin, D2, Ireland

b Nexalus Labs, Water St, Loughanalla, Castlepollard, Co Westmeath, Ireland

c CONNECT, Dunlop Oriel House, Trinity College Dublin, Ireland.

Abstract

This case study presents an experimental investigation of closed loop liquid cooling of CPUs. The cooling system implements a pumped liquid supplying a cold plate fixed to the emulated CPU heat source and a remote Air-Cooled Heat Exchanger (ACHE). The Hybrid Liquid-Air (HLA) system employs 25% v/v Arteco-Freecor with distilled water solution as the primary coolant, and air as the secondary coolant. The thermal-hydraulic performance of the system is characterized over a range of operating conditions, including CPU heat loads (12 W/cm^2 - 72 W/cm^2), water volumetric flow rates (0.5 L/min - 8 L/min), and air velocities (1m/s – 3m/s). In order to provide a performance frame of reference, evaluations were also conducted with a conventional air-based fan-fin CPU cooling system.

The results indicate that the HLA system significantly outperforms the traditional air-cooled fan-fin system, with much lower CPU temperatures and feasible power levels exceeding 300 W. Additionally, the study highlights the importance of the liquid flow rate and fan speed on the overall cooling performance. Finally, an optimal combination of liquid and air flow rates is identified for minimum energy consumption. Overall, these findings support the conclusion that conventional air-cooling is reaching its theoretical limit and closed loop liquid cooling has the potential to keep pace with the escalating power levels of modern CPUs. However, the additional energy cost of the cooling system should be carefully managed in order to maximize overall system Coefficient of Performance.

Keywords:
CPU Cooling, Liquid Cooling, Hybrid Cooling, Air-Cooling, Energy Optimisation

1. Introduction

The demand for high-performance computing has led CPU manufacturers to develop increasingly powerful processors. These CPUs continue to consume more power and generate higher heat fluxes, averaging about 70 W/cm^2 with Thermal Design Powers (TDPs) exceeding 150 W [1]. Intel's i9 7090XE, 56 Xeon Platinum 8470Q (formerly Sapphire Rapids), and 56 Core Xeon Platinum 9282, among others, exemplify this trend with power outputs of 280 W (overclocked), 350 W, and 400 W respectively [2]. Managing the heat generated by these CPUs has become a significant and important challenge, since effective cooling solutions are required to prevent thermal throttling or damage to the CPU.

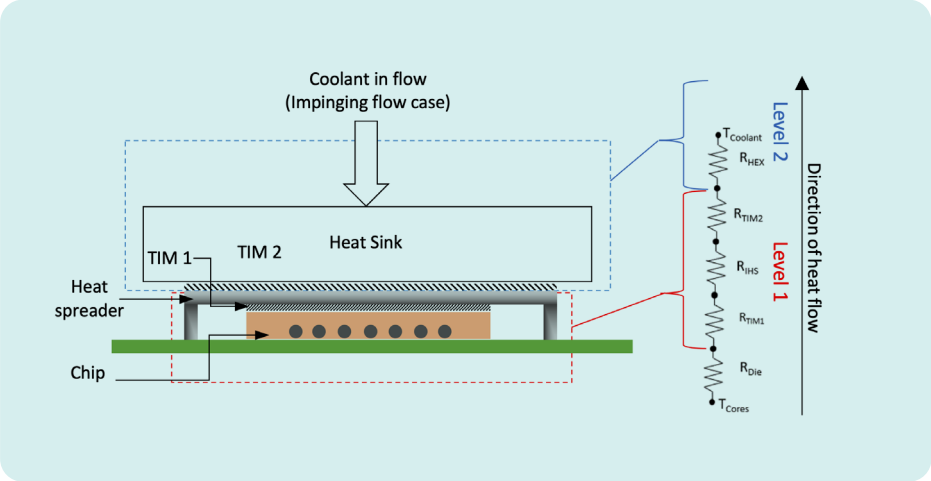

Fig 1. Thermal resistance network of a CPU Package.

A simplified thermal resistance network of a CPU packaging is shown in Figure 1. The core to coolant thermal resistance consists of five resistances in series, divided into two levels: level one includes the resistance R_{die} due to the conduction through the semiconductor die, the resistance R_{TIM1} across the Thermal Interface Material (TIM) between die and heat spreader, and the spreading resistance R_{IHS} within the Integrated Heat Spreader (IHS); level two consists of resistance R_{TIM2} across the TIM between IHS and base of the heat exchanger, and the resistance R_{HEX} associated with the heat exchanging technology to the coolant. Among these, the first-level resistance is challenging to regulate, as it falls under the jurisdiction of the chip manufacturer. Nevertheless, the resistances at the second level, in particular R_{HEX}, can be engineered by electronic packagers to reduce R_{Total}.

The majority of CPUs, whether in PCs or servers, are directly cooled by fan-fin heat sinks. These heat sinks can have duct (parallel) or impinging (normal) air flow configurations [3]. Desktop PCs typically use impingement cooling, while servers use ducted cooling. As CPUs push the boundaries of performance and power consumption, the limitations of air cooling become increasingly evident. The cooling threshold for air cooled fan-fin heat sinks can be as low as 40 W/cm^2 [4]. Despite efforts to optimize heat sink designs and enhance fan performance, there exists a critical point where air cooling can no longer effectively cope with the escalating heat output of high-performance CPUs. This limitation can push power-hungry CPUs to their thermal limits, putting system performance, stability, and reliability at risk. The limitation of air cooling becomes more pronounced when considering the non-uniform heat distribution of modern CPUs [5]. In addition to the heat transfer limitations of fan-fin air cooling, other critical factors like reliability, leakage current, and harsh operating conditions (especially at elevated temperatures), have emerged as strong motivators for the adoption of liquid cooling technology. Despite this, adoption rates have been slow. While air cooling remains a practical and cost-effective solution for many everyday computing needs, it is important to recognize its limitations in the face of ever-evolving CPU thermal challenges. The pursuit of cutting-edge performance demands high performance thermal management technologies, and liquid cooling is emerging as an obvious choice for modern high-powered CPUs.

Motivated by the superior thermophysical properties of liquids, single-phase liquid cooling first emerged in the mid-1960s as a means to cool large mainframe computers. However, it was not until 1981 that Tuckerman and Pease [6] introduced the concept of direct-to-chip liquid cooling, using microchannels to enable heat flux removal rates exceeding 790 W/cm^2. However, with the advent and uptake of Complementary Metal Oxide-Semiconductor (CMOS) transistors in the 1990s, the computing industry shifted its attention away from liquid cooling. This was primarily due to the reduced heat generated by CMOS technology, which made air cooling a feasible and cost-effective solution. Nonetheless, in recent years, there has been a renewed interest in liquid cooling technologies due to the substantial increase in power densities of transistors and the realization that air cooling is approaching its theoretical limit.

Microchannel-based liquid cold plates have been extensively researched as a high potential cooling solution for high heat flux electronics, primarily owing to the fact that smaller hydraulic diameter significantly increases the convective heat transfer coefficient, whilst the interconnected microchannels create high surface area density [7]. With a relatively slow start, there has been renewed interest in microchannel heat sinks for cooling of electronics since the early years of this century, and research efforts have centred around improving heat transfer capabilities while managing pressure drop penalties. Various concepts have been explored, including considerations of channel geometry, flow obstructions, wavy and zig-zag channels, and introducing periodic pinching and channel roughness [8–11]. This is not an exhaustive inventory of designs conceived to enhance the thermal-hydraulics of microchannel heat exchangers, yet serves as a testament to the availability of diverse design alternatives for engineering such technologies. Research highlights include the studies of Colgan et al. [12] who achieved high overall heat transfer coefficients >200 kW/m^{2}K, and reported heat flux dissipation >1000 W/cm^2 [13]. Nevertheless, the utilization of reduced channel sizes in microchannel heat sinks results in the trade-off of higher pressure drop and hydraulic power requirements [14]. Also, temperature non-uniformities can arise in microchannel heat exchangers due to the sensible heating of the fluid along the channel length, nor are they well-suited for managing hotspots [14]. The presence of temperature gradients in electronics has long been known as detrimental to their reliability due to the thermal stresses induced by differential thermal expansion [15].

Alternatively, liquid microjet array impingement is a complimentary approach to microchannels as they achieve remarkably high heat transfer rates without the need for surface area enhancements. For example, in Michna et al. [16] a microjet array dissipated heat flux of 1100 W/cm^2, albeit from a very small target area. Similar to microchannel flows, the convective heat transfer coefficient increases as the size of the microjets decreases [17,18]. Despite being relatively less explored in the available literature, microjets are emerging as a formidable contender for high heat flux cooling applications due to their ability to generate substantially higher convective heat transfer coefficients compared to traditional channel-based designs, and generally with lower pressure drop penalties [19]. The comparatively high convective heat transfer coefficients are attributed to the high velocity of the impinging jets, which results in a high stagnation heat transfer coefficient at the heated surface. By dispersing coolant across the entire heated surface, jet arrays also have the potential to greatly reduce undesired surface temperature gradients. Furthermore, microjet arrays can be designed to provide targeting of hotspots to improve surface temperature uniformity [19]. However, the heat transfer between neighbouring jets in the array tends to degrade rapidly, presenting an engineering challenge in achieving high cooling effectiveness over large surface areas at reasonable volumetric flow rates [20].

Hybrid microjet-microchannel systems have also been studied. Here, the liquid is discharged through the perforated orifice plate as multiple microjets into a microchannel feature on the cold plate, thus capitalising on the merits of each technology to achieve the high cooling intensity whilst mitigating high temperature gradients and hydraulic penalties. For example, Robinson et al. [21], achieved effective conductance levels >400 kW/m^2 K, at a moderate pressure drop (< 100 kPa).

1. Introduction

In a closed loop configuration, ensuring optimal performance relies not only on the cold plate heat sink but also on the integration of a reliable and compact pump to facilitate fluid circulation within the loop. Additionally, a liquid-air heat exchanger is required to dissipate the captured heat from the cold plate. Furthermore, improved heat transfer provided by the liquid cold plate comes at the price of the pressure drop and associated pumping power. This, and any additional fan power requirements, must be considered in the context of the overall energy consumption of the cooling system.

From a system integration perspective, existing research has primarily focused on the CPU alone, whereas the full source-to-sink thermal-hydraulics has been largely overlooked. Some studies aimed to address this gap by investigating the overall thermal resistance of liquid cooling systems, including the remote heat exchanger and the interplay of all system components. A compact CPU cooling system was studied by Bintoro et al. [22], comprising a single water jet impingement device and mini-channel remote heat exchanger, achieving an overall source-to-sink thermal resistance of approximately 0.35 K/W. Similarly, Liu et al. [23] and Chang et al. [24] used microjets and microchannel heat sinks respectively, coupled with a remote heat exchanger to achieve overall thermal resistances of about 0.17 K/W–0.23 K/W. Efforts by Whelan et al. [25] to design and develop a full closed loop CPU liquid cooling system achieved an overall thermal resistance of 0.18 K/W, with a low pumping power of 1.5 W.

Considering the above discussion, the current study is motivated by the recognition that high-powered CPUs will necessitate cooling solutions that surpass the inevitable limitations of conventional fan-fin air cooling. Single-phase liquid cooling emerges as a very high potential option from technical and practical standpoints. However, existing literature highlights knowledge gaps concerning the system-level assessment of closed loop liquid-based cooling technology. These gaps encompass full system-level thermal-hydraulic performance characteristics across a range of operating conditions, as well as energy management considerations, particularly in terms of characterizing and optimizing the energy penalty associated with the full cooling system. Thus, the objective of this research is to conduct a comprehensive case study on CPU closed loop liquid cooling, encompassing the thermal-hydraulics from the heat source to the heat sink. By performing a thorough parametric study across a sufficient range of system parameters, the research aims to fill gaps in the existing literature by gaining a deeper understanding of the parameter sensitivity, and its overall influence on performance for end-to-end systems. Since this investigation compares the performance characteristics with traditional fan-fin technology, it also presents electronics packagers with missing information, as well as a methodology and characterization procedure, for making practical decisions when evaluating the limitations of conventional air cooling. Finally, the investigation extends to optimizing the energy cost associated with closed loop liquid cooling.

2. Experimental Apparatus and Data Reduction

2.1 Thermal Test Vehicle with Liquid-Cooled Cold plate and Remote Air-Cooled Heat Exchanger

One key aim of this investigation is to evaluate the thermal-hydraulic performance characteristics of a pumped Hybrid Liquid-Air (HLA) system consisting of a cold plate and remote Air-Cooled Heat Exchanger (ACHE). To achieve this, an instrumented flow loop facility that emulates the practical implementation of this type of cooling solution was constructed. As Figure 1 indicates, the flow loop has three regions of interest – the Thermal Test Vehicle (TTV), the primary (liquid) coolant loop, and the ACHE. As will be discussed, the TTV is comprised of a variable heat source which is cooled by a liquid cold plate. The present study employs a coolant mixture comprising 25% v/v Arteco-Freecor FTC® and 75% v/v distilled water as the primary coolant. This coolant mixture was selected to provide freeze protection and inhibit corrosion and microbial growth compared to water [26].

The primary loop includes a SAER CF gear pump with a AC10 series Parker Hannifin single phase inverter drive that controls the pump speed to the desired liquid flow rate. The flow rate is measured by a non-invasive UF25B Cynergy ultrasonic flow meter with uncertainty of ± 3 % of reading. Pressure drop measurements are taken across both the cold plate and the ACHE with a 26PCDFA6D Honeywell differential pressure transducer with uncertainty of ±0.25%. Similarly, temperature differentials are taken across both the cold plate and the ACHE with T-Type thermocouples, which were inserted into the flow via compression fittings. The temperature probes were calibrated against a platinum RTD reference probe to an uncertainty of a ±0.2 °C.

A Thermal Test Vehicle (TTV) was constructed to represent a CPU heat source. As depicted in Figure 3 (a), the TTV was machined from single block of ultra-pure OFE copper, with a lower neck region of A_{die}=20 x 20 mm^2 square cross section, roughly the area of higher powered (>150 W) commercial CPU dies (e.g. Intel i9 7980XE), which tend to be marginally larger than their lower-powered counterparts. Three OMEGALUX® CS series stainless steel sheath cartridge heaters of 6.35 mm diameter that could supply a maximum total power of 100 W each were inserted into the bottom of the 60 mm long lower neck section to provide heat to the upper 40mm x 40 mm cross section and 2.5mm thick overhang section, which represents a CPU Integrated Heat Spreader (IHS). The TTV section is then housed in a 3D printed case lined with insulation to minimized heat losses and provide a mounting fixture for the attached liquid-cooled cold plate.

The measurement of the heat flux from the lower neck to the IHS, as well as the temperature of the tip of the neck region, here representing the CPU die temperature, is facilitated by three 1.0 mm diameter T-Type thermocouples spaced at equal intervals along its length and inserted into drilled holes to a depth of 10 mm. As Figure 3 (b) illustrates, this allows the determination of the temperature gradient at the base of the IHS, dT/dx, such that the associated heat transfer to the IHS section can be determined from,

$$Q = -k_{Cu} A_{die} \frac{dT}{dx} \tag{1}$$

The temperature of the neck-IHS interface, T_4 (hereinafter T_{die}), represents the temperature of the CPU die and is determined from linear extrapolation,

$$T_{die} = T_3 - \left(\frac{Q}{A_{die}}\right)\left(\frac{\Delta x_3}{k_{Cu}}\right) \tag{2}$$

The IHS above the heat source was cooled by a commercial cold plate (Nexalus Enflux™). As Figure 2 indicates, this cold plate consists of a plastic manifold that forces the liquid through an array of 0.7 mm orifices, creating an array of high velocity microjets that impinge on a nickel-plated copper baseplate. The cold plate was mounted to the TTV via four bolts tightened with a torque wrench to ensure correct and even pressure. A layer of Nexalus Hydronex™ thermal grease is placed between the IHS and the cold plate to reduce contact resistance, typically referred to as TIM2.

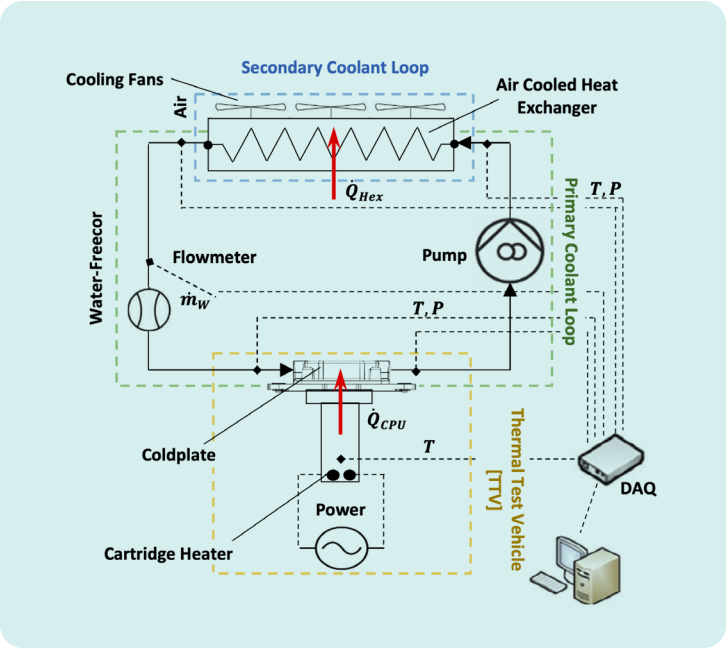

Fig 2. Wireframe diagram of the Liquid Cooling System using Air Cooled Heat Exchanger System

With the heat transfer to the IHS measured, the overall die-to-coolant thermal resistance can be estimated as.

$$R_{d-c} = \frac{\Delta T_{cp}}{Q} \quad (3)$$

As suggested by Webb [27], the temperature difference across the cold plate, ΔT_cp, is calculated as the Log-Mean Temperature Difference (LMTD),

$$\Delta T_{cp} = \frac{T_{out} - T_{in}}{\ln\left(\frac{T_{die} - T_{in}}{T_{die} - T_{out}}\right)} \quad (4)$$

It must be noted that overall thermal resistance includes the resistance contributions of the IHS, TIM2 and the cold plate heat exchanger. In a commercial CPU package, additional thermal resistances associated with TIM1 needs to be included for accurate estimation of the die temperature.

Consistent with modern liquid cooling technology currently deployed in high-performance desktop/workstation computing systems, the heat rejection to the air was facilitated via a remote air-liquid heat exchanger. In this investigation, the Nexalus NXQ 360 fan-cooled liquid-air flat tube heat exchanger of outer dimensions 400 mm x 130 mm x 60 mm, fitted with three 120 mm high static pressure fans, was used to transfer the heat from the liquid to the air. The heat exchanger is equipped with 16 flat copper tubes, each 2mm wide, spanning its entire length. The copper fins of the radiator are non-louvred and have a density of 18 fins per inch (FPI). Thermocouples were positioned in the entrance and exit air streams in order to evaluate the temperature rise across the air-side of the heat exchanger. These measurements, along with those of the water-side allow for the calculation of the overall thermal resistance of the heat exchanger,

$$R_{hex} = \frac{1}{(UA)_{hex}} = \frac{\Delta T_{hex}}{Q} \tag{5}$$

where ΔT_{hex} is the LMTD of the crossflow liquid-air heat exchanger and is calculated by,

$$\Delta T_{ACHE} = \frac{\left(T_{hex,out} - T_{air,in}\right) - \left(T_{hex,in} - T_{air,out}\right)}{\ln\left(\frac{\left(T_{hex,out} - T_{air,in}\right)}{\left(T_{hex,in} - T_{air,out}\right)}\right)} \tag{6}$$

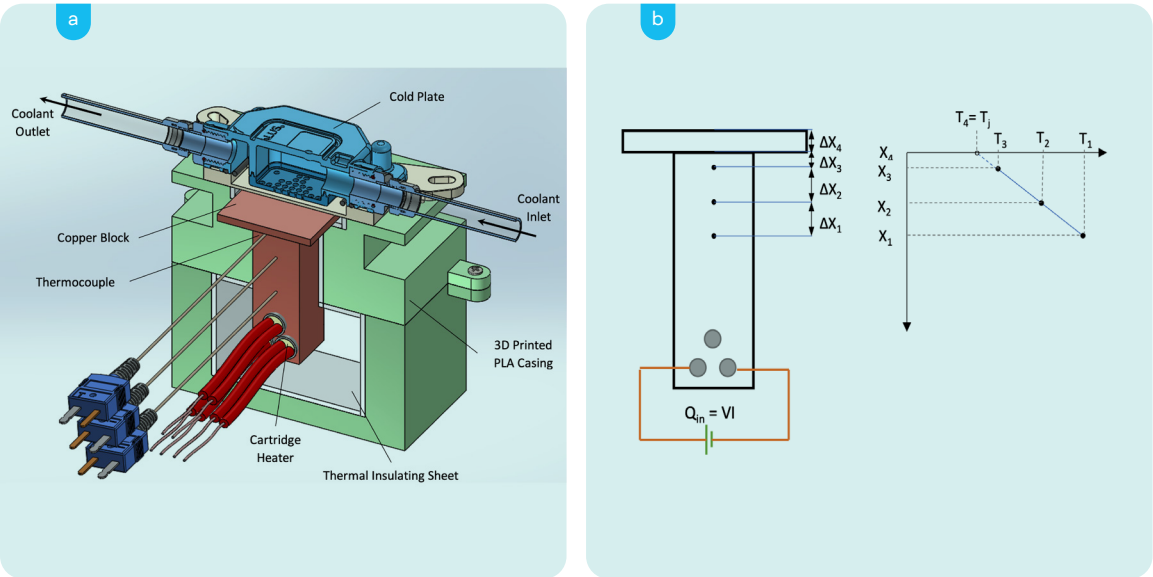

Fig 3. Schematic of the Thermal Test Vehicle (TTV) with mounted liquid-cooled cold plate, and (b) illustration of heated meter bar and the resulting linear Temperature vs. Distance plot.

The fan speed was controlled by varying the supply voltage with a desktop DC power supply over the range of 6V-14V. The penalty associated with the power required for the liquid-side of the cooling system requires determination of the hydraulic pumping power across the cold plate and ACHE,

$$\dot{W}_{pumping} = \dot{V}(\Delta P_{cp} + \Delta P_{ACHE})$$

(7)

where V̇ is the liquid volumetric flow rate and ΔP is the associated pressure drop. It is important to note that the pressure drops arising from the valves, T-joints, flow meters are not considered here as these would be absent in a real-life application. The total pumping power along with the total fan power,

$$\dot{W}_{fans} = IV$$

(8)

allows for the estimation of the overall system Coefficient of Performance (COP),

$$COP_{ACHE} = \frac{Q}{\dot{W}_{pumping} + \dot{W}_{fans}}$$

(9)

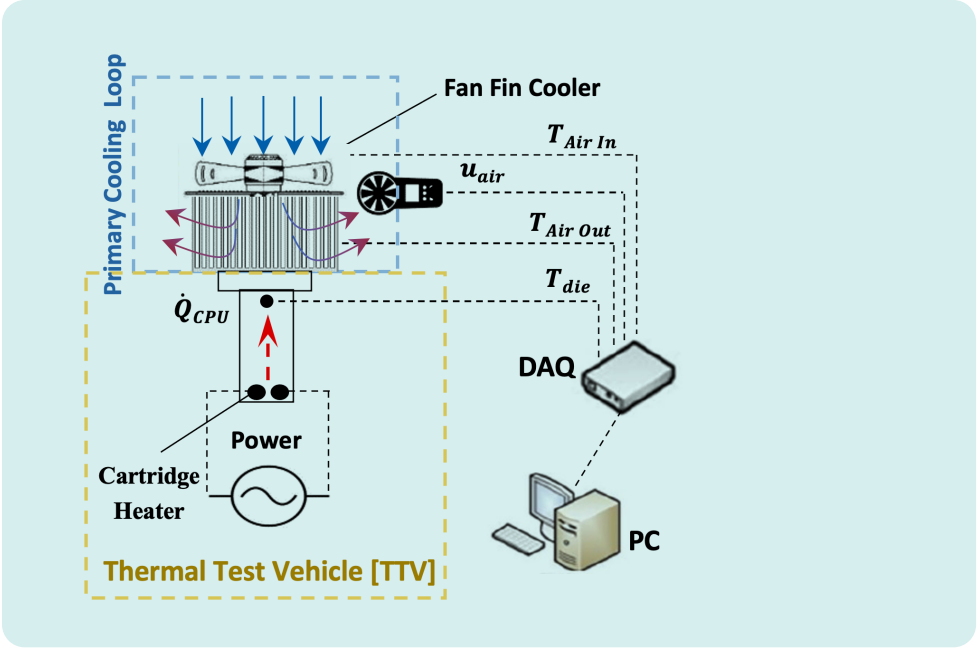

Fig 4. Wireframe diagram of the impingement flow Fan-Fin Air Cooler (FFAC) system.

2.2 Fan-Fin CPU Cooler

In order to establish a performance frame of reference for comparison with the liquid-cooled system described above, commercial off-the-shelf Fan-Fin Air Cooler (FFAC) was also characterized. The FFAC used was typical of most commercial impingement type heat sinks, with aluminium fins measuring 95mm x 98mm x 70mm and equipped with a 92mm fan. The heat sink was affixed to the Integrated Heat Spreader (IHS) of the TTV via mechanical fastening, with the same layer of thermal grease (TIM2) placed between the IHS and the aluminium heatsink base, as depicted in Figure 4. The heat sink had a total of 35 fins, each of 0.6 mm thickness and 34.5 mm height, with 1.8 mm spacing between them.

In characterizing the thermal performance of the FFAC, the source-to-sink thermal resistance (R_{FFAC}) is calculated as,

$$R_{FFAC} = \frac{\Delta T_{FFAC}}{Q} \tag{10}$$

where ΔT_{FFAC} is the LMTD of the cold plate heatsink and is calculated by

$$\Delta T_{cp} = \frac{T_{air,out} - T_{air,in}}{\ln\left(\frac{T_{die} - T_{air,in}}{T_{die} - T_{air,out}}\right)} \tag{11}$$

By utilizing a benchtop DC power supply, the fan speed was regulated, and the total power was computed similar to Eq. (8), such that the system's Coefficient of Performance (COP) can be determined from,

$$COP_{FFAC} = \frac{Q}{\dot{W}_{fans}} \tag{12}$$

2.3 Uncertainty Analysis

The uncertainties associated with the primary measurement instruments are provided in Table 1. To determine the propagated uncertainty of calculated quantities, standard methodologies outlined by Taylor [28] were followed, assuming that the uncertainties were random and independent. In cases where a calculated quantity Z relied on measured quantities X, Y, and possibly others, the uncertainty δZ was determined using the following prescribed procedures.

$$\delta Z = |Z| \sqrt{\left(\frac{\delta X}{|X|}\right)^2 + \left(\frac{\delta Y}{|Y|}\right)^2 + \cdots} \quad (13)$$

The uncertainty associated with the rate of heat transfer through the heated meter bar system requires some special attention, as it depends not only on multiple temperature and spatial errors, but also the regression fitting to these to determine *dT/dx*. Similarly, this error propagates to the estimation of T_{die}, which must also be quantified. To account for this, the Monte-Carlo calculation method first described in Kempers et al. [29] and later detailed in Kempers and Robinson [30] was implemented.

Measurement	Instrument	Absolute Uncertainty
Temperature	T-type Thermocouple	± 0.2 °C
Flowrate	Ultrasonic Flowmeter	± 3 %
Distance	Vernier Calliper	± 0.0005 m
Pressure Drop	Pressure Sensors	± 0.25 %
Air Velocity	Vane Anemometer	± 2 %

Table 1: Uncertainties of measured quantities

3. Results and Discussion

3.1 Thermal Characterisation

3.1.1 Fan-Fin Air Cooling System

Figure 5 shows that, for the FFAC system, as the heat flux supplied from the TTV increases from 12.25 W/cm^2 to 72 W/cm^2 the net thermal resistance remains relatively constant within its absolute uncertainties. Increasing the air velocity across the heat sink lowers its thermal resistance from about 0.32 K/W to 0.27 K/W, indicating marginally improved heat transfer. Figure 6 shows the expected cooling performance of the FFAC system under various heat loads for an ambient temperature of 35°C, the standard room temperature recommended by Intel [31]. Under extreme loads, the average die temperatures can reach as high as 80–85°C, and this is regarded as a 'safe' operating temperature limit [32]. The throttling temperature limit may differ depending on the CPU type, typically falling within the range of 100–105°C [33]. In this study, the safe and throttling temperature limits of 83°C and 102°C respectively, are chosen when comparing cooling technologies. To predict the die temperature closer to what would be expected in a CPU package, an approximation for the TIM1 thermal resistance, using the R_{TIM1} = 0.1 K-cm2/W, was included [4].

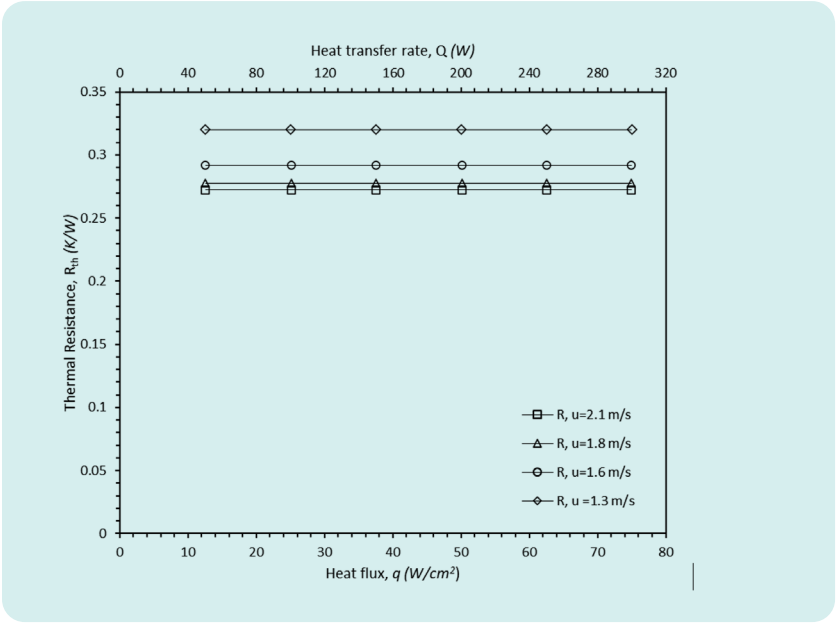

Fig 5. Thermal resistance of fan-fin based air cooling system with varied heat loads; $T_{ambient}$ = 20°C.

For a normal fan speed of 1.8 m/s, the safe die temperature limit of 83 °C was breached at approximately 35 W/cm^2, which is close to that reported in the literature [4]. The temperature at the maximum heat flux studied (72 W/cm2) was at least 53K higher than the safe limit, and even well above the throttling temperature of 102°C [33]. Consistent with the thermal resistance behaviour, there is a noticeable yet marginal effect of air velocity on the die temperature which diminishes with increased velocity. Considering the diminishing gain in the heat transfer, together with the escalation in fan noise and power consumption, it is not worthwhile to increase the air speed beyond approximately 1.8 m/s.

As Figure 7 indicates, increasing the heat flux results in an increase in the Coefficient of Performance (COP) of the system, ranging approximately from ~15 to 90 for settings that achieve lower than the throttling temperature. This result is expected since, despite the increase in heat transfer rate, the fan power required to cool CPU was kept constant. However, as Figure 7 shows, higher fan speeds lead to a significant decrease in the COP. This occurs because, even though the heat flux remains constant, the fan power consumption nearly doubles when the air velocity changes from 1.3 m/s to 2.1 m/s, resulting in the COP being halved. Upon closer examination of the temperature and COP plots, it was observed that surpassing an air velocity of 1.8 m/s does not provide substantial cooling benefits. Instead, it leads to increased power consumption, lower COP values, and heightens fan noise. This highlights a key point which has not been addressed sufficiently in the literature: a balance needs to be struck between cooling requirements and the associated penalties, like energy consumption and noise. Each aspect must be considered holistically in the development of a cooling system.

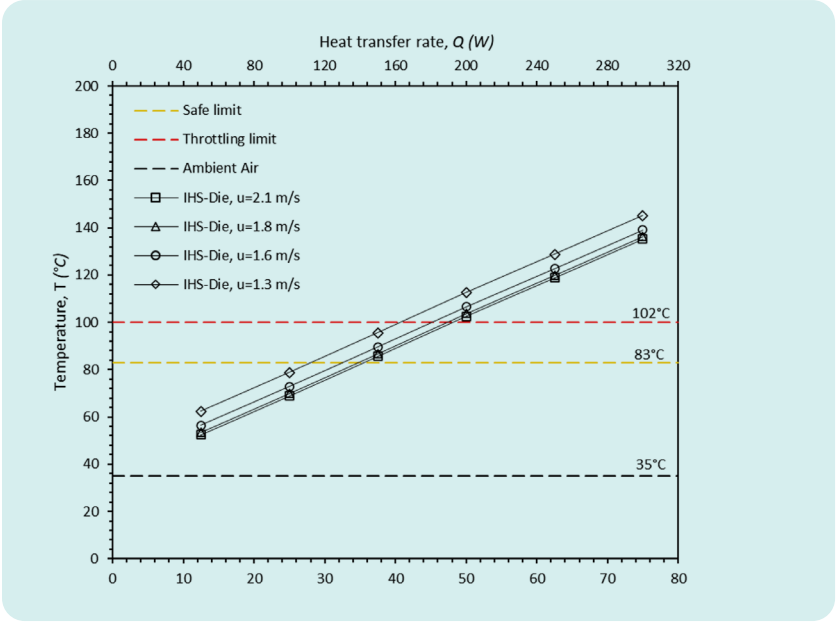

Fig 6. Expected cooling performance of a standard Fan & Fin Air Cooler at maximum fan speed with varying heat loads for $T_{ambient}$ =35°C and R_{TIM1} = 0.1 K-cm2/W

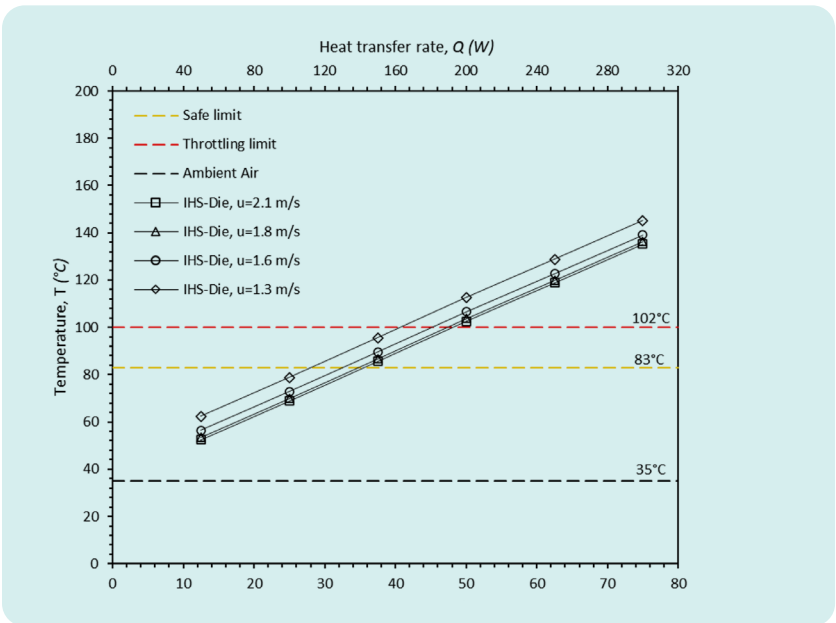

Fig 7. Power Consumption & Coefficient of Performance (COP) profile of the fan-fin heat sink; $T_{ambient}$ = 20°C.

3.1.2 Hybrid Liquid-Air (HLA) Cooling System

Figure 8 shows that for the selected nominal settings (max fan speed & V̇=4 L/min), the thermal resistance of the cold plate and ACHE remains almost constant, as the heat flux increases from 12.25 W/cm² to 72 W/cm². It is worth noting that the thermal resistance of the cold plate is roughly ten times that of the ACHE, which supports the focus on cold plate research in the literature. Notably, the net thermal resistance of the liquid cooling system is about 0.11. K/W, which three times lower than the FFAC, which is significant.

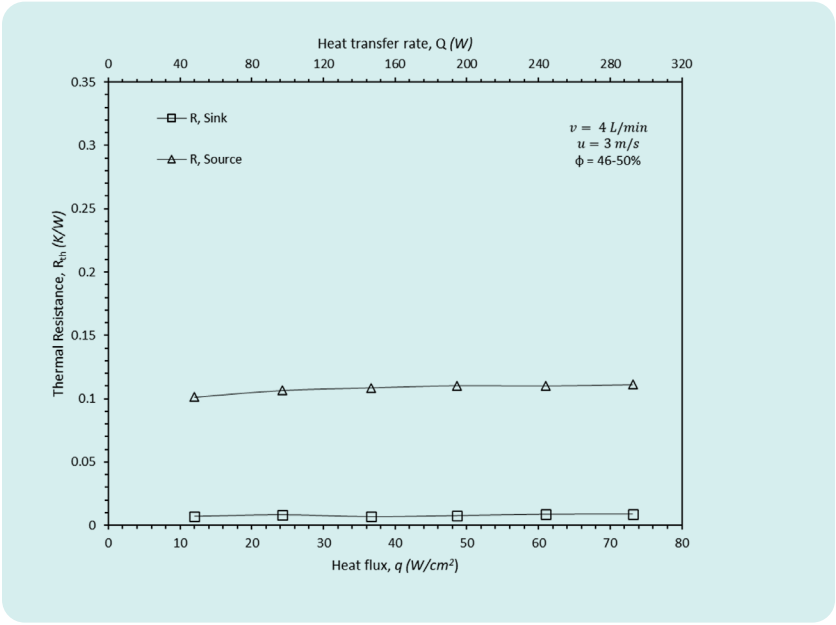

Fig 8. Thermal Resistance of HLA cooling system with varied heat loads: $T_{ambient}$ = 20°C.

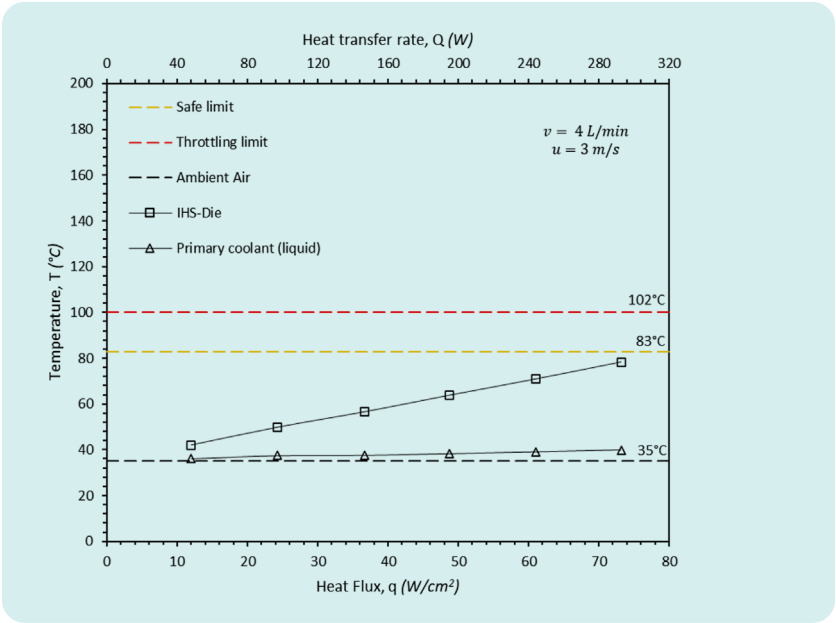

Fig 9. Expected cooling performance of HLA cooling system at varying heat loads for a nominal flow rate of 4 L/min and fan speed of 3m/s at $T_{ambient}$ = 35 °C and R_{TIM1} = 0.1 K-cm2/W.

The expected die temperature with increasing heat flux are plotted in Figure 9, again considering $T_{ambient}$ = 35°C and approximate TIM1 thermal resistance. As the primary and secondary coolant temperatures experience only slight increases over the range of considered heat fluxes, T_{die} increases linearly with the heat transfer rate, attributable to the constant source-to-sink thermal resistance. The HLA system maintains temperatures well within the safe operating limits for the full range of high heat fluxes tested, and throttling is not expected until power levels of close to 450 W.
When comparing the performance of the HLA and FFAC systems, it was found that the HLA system is capable of cooling the die by at least 60 K more effectively than the fan-based system for the highest tested heat flux of 72 W/cm2.

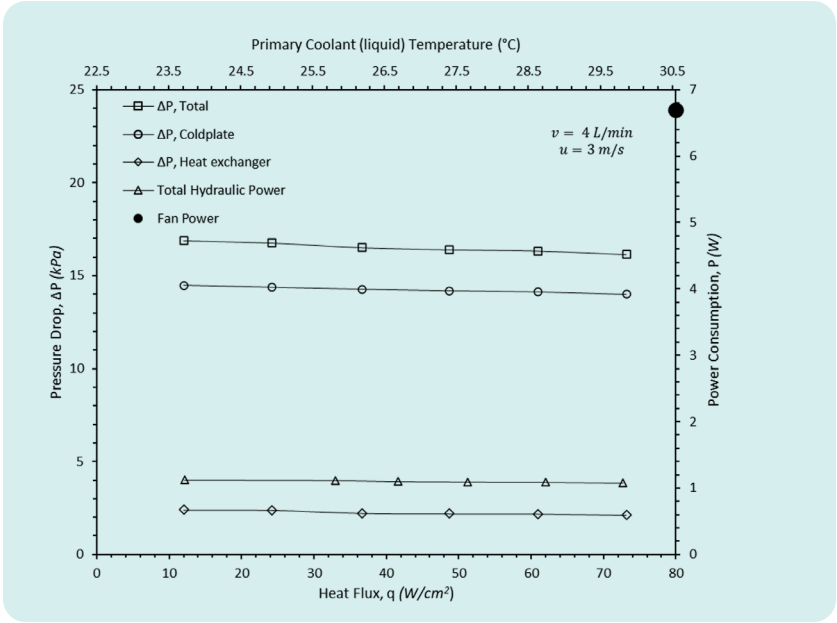

Fig 10. Pressure drop curves of the HLA system for varying heat fluxes at constant coolant flow rates: $T_{ambient}$ = 20°C.

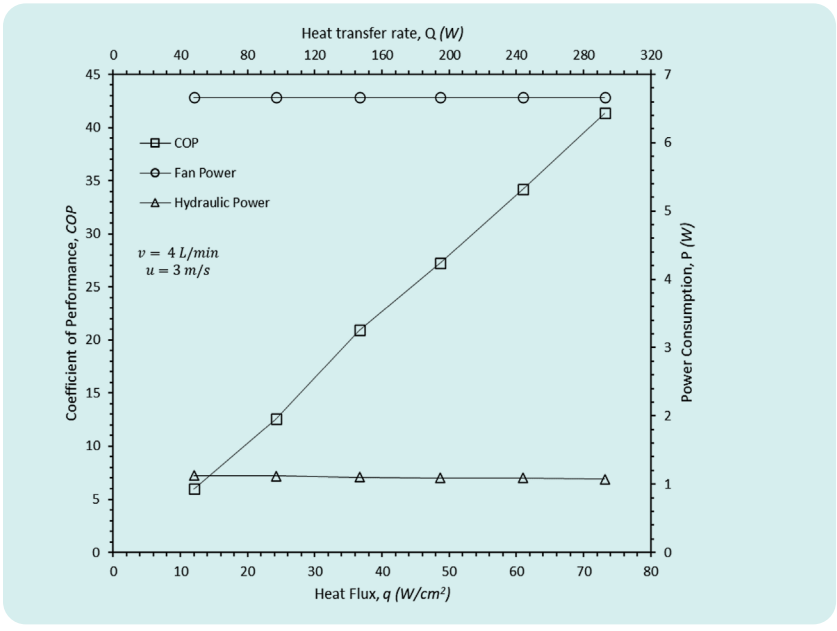

Fig 11. Power Consumption & Coefficient of Performance (COP) profile of the HLA system; $T_{ambient}$ = 20°C.

Figure 10 plots the HLA system pressure drops at varying heat load for the nominal flow rate of 4 L/min across both the cold plate and ACHE. The cold plate demonstrates an expectedly higher pressure drop compared to ACHE (approximately 16 kPa compared to 2 kPa), though these combine to a pumping power of only 1.1 W. It is worth noting that fan power at the ACHE for this maximum air velocity is nearly six times greater than the hydraulic power necessary to circulate the liquid in the loop, as shown in Figure 11. Therefore, previous research efforts that focused on enhancing cold plate heat transfer capabilities whilst managing the pressure drop is justified.

As seen in Figure 11, the COP of the system is observed to rise linearly, as expected, from approximately 5 to 36 over the range of TTV powers tested. A slight decrease in hydraulic power is observed due to the lower pressure drop in the liquid loop caused by a slightly lower coolant viscosity at the higher operating temperatures, as seen in Figure 10. Compared to the FFAC, the COP of the HLA system is approximately 2.5X lower for settings that achieve lower than the throttling temperature. However, this penalty must be considered in the context of the benefits associated with the lower operating CPU temperatures and extended power range afforded by the HLA system.

3.2 Liquid Flow Characterisation – HLA System

Figure 12 shows the variation of thermal resistance with liquid flow rate for a fixed maximum heat flux and fan speed. The plot shows that as the liquid flow rate increased from 0.5 L/min to 8 L/min, the thermal resistance decreased, though its effect is much pronounced for the cold plate compared to the ACHE. However, beyond around 4 L/min, the flow rate does not have significant effect on the thermal resistance. This is also noticeable from the temperature plots in Figure 13, whereby any increase in coolant flow rate beyond 4 L/min has diminishing influence on the die temperature, which are all predicted to be below the safe operating temperature for this maximum heat flux rest case. This diminishing return with liquid flow rate is expected since (i) jet array impingement heat transfer increases as Nu ∝ Red^n, with n~0.5 [17], and (ii) as the convective resistance decreases, those of the IHS and TIM2 become increasingly dominant.

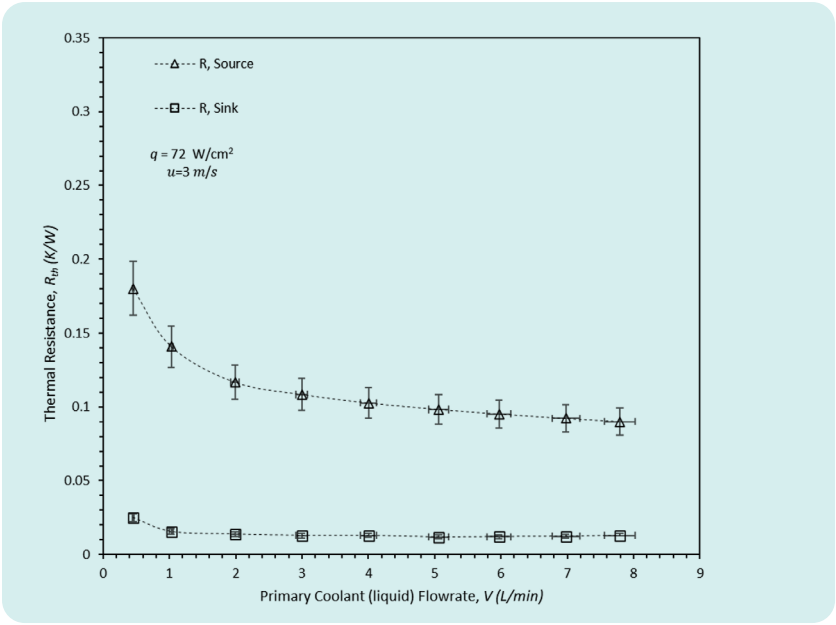

Fig 12. Thermal Resistance of HLA system for varying liquid flow rates measured at constant heat flux and air velocity for $T_{ambient}$ = 20°C.

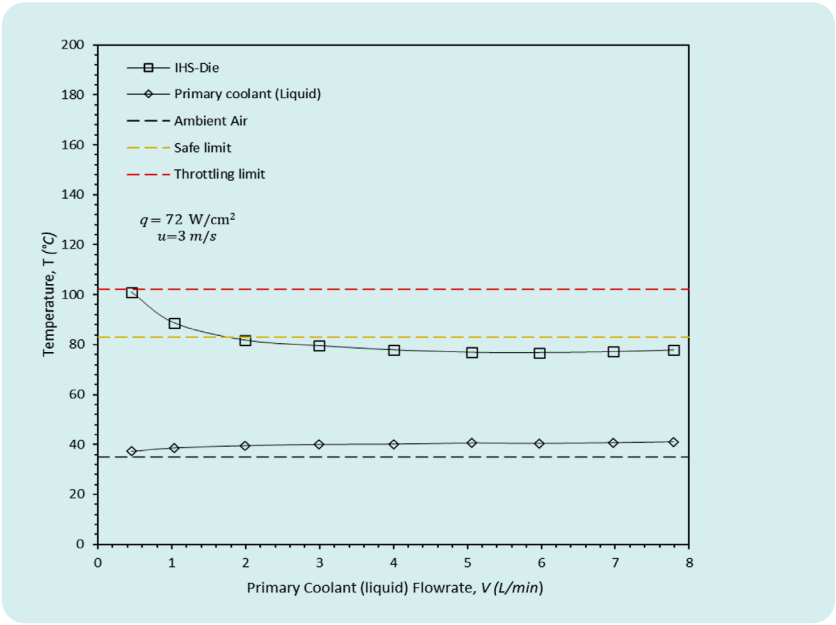

Fig 13. Expected Cooling Performance curves of HLA system for varying liquid flow rates at constant heat flux and air velocity at $T_{ambient}$ = 35 °C and R_{TIM1} = 0.1 K-cm2/W.

However, as shown in Figure 14, the pressure drop across the cold plate and ACHE increases with the expected parabolic profile, with the former dominating the total system pressure drop and thus hydraulic power. As such, the diminishing gains in heat transfer associated with flow rates beyond about 4 L/min has associated with it increasingly severe hydraulic penalties. Again, this illustrates that, when considered as a full system, careful design consideration must be made to ensure adequate performance with managed hydraulic penalties. This is particularly important considering the impact on the overall energy dynamics of the HLA system. This is displayed in Figure 14, where increasing flow rates from 4 L/min to 8 L/min leads to about a 3.5-fold increase in the hydraulic power, which negatively affects the system COP, as shown in Figure 15, with little improvement in the cooling performance (Figure 12). However, this is moderated to some degree by the fan power (~6.6 W), which is notably higher than the hydraulic power for liquid flow rates up to 4 L/min. This causes the COP to level-off at lower liquid flow rates, and to a lower value for decreasing heat input levels (see Figure 15).

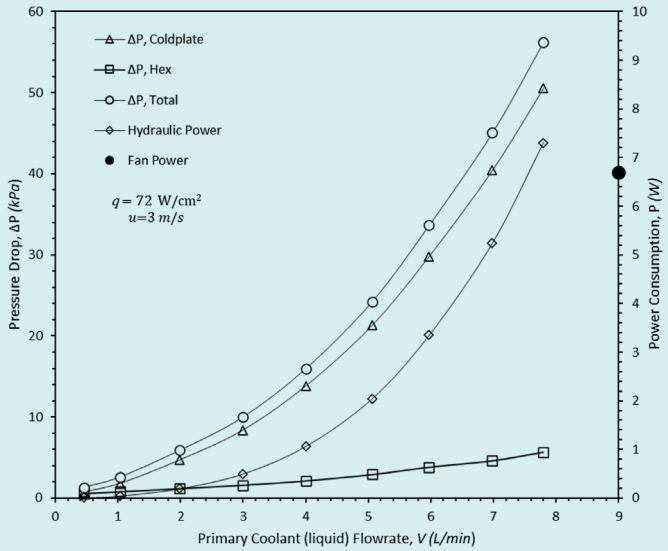

Fig 14. Pressure drop plots of HLA for varying liquid flow rates at constant heat load: $T_{ambient}$ = 20°C.

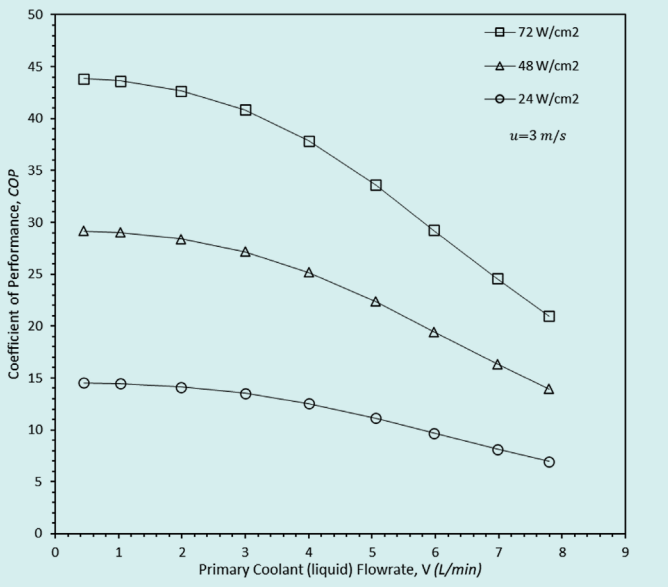

Fig 15. COP plots of HLA system for three different heat fluxes with varying liquid flow rates measured at constant air velocity: $T_{ambient}$ = 20°C.

3.3 Fluid Characterisation – Air side

Keeping the heat flux constant (72 W/cm^2) and liquid flow rate at its nominal 4 L/min, the influence of the air velocity (fan speed) was varied between 1 to 3 m/s to investigate its overall influence on system performance. As evident from Figure 16, increasing the air velocity has much less significant influence on the ACHE thermal resistance compared with the effect the liquid flow rate has on the cold plate resistance (Figure 12). It is thus unsurprising that the predicted die temperature is nearly constant with increasing fan speed, as seen in Figure 17. The same cannot be said about the power consumption and related COP, which depicted in Figure 18, shows the system COP decreasing steeply with increased fan speed. This is due to the increased power consumption of the fan, which increases about 5.5-fold. Notably, Figure 18 shows that the fan speed can be almost as low as the pumping power at the lowest fan speed setting, and this occurs with only a minor drop in the die temperature (Figure 17).

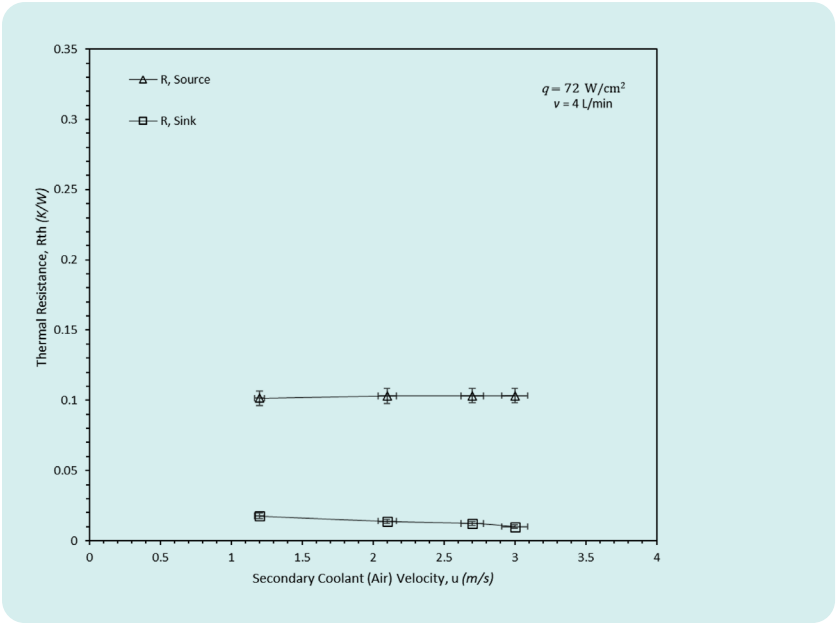

Fig 16. Thermal Resistance of HLA system for varying air velocities measured at constant heat flux and liquid flow rate for $T_{ambient}$ = 20°C.

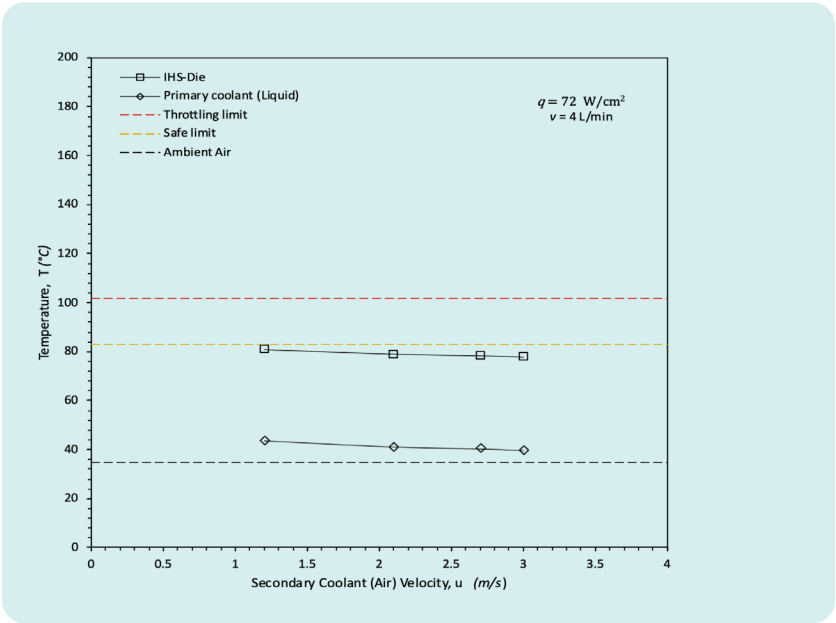

Fig 17. Temperature profile of the HLA system for varying air velocities measured at constant heat flux and liquid flow rate at $T_{ambient}$ = 35°C and R_{TIM1} = 0.1 K-cm2/W.

This investigation of closed loop liquid cooling on high-powered CPUs showed significantly higher cooling capacity than the traditional FFAC system. For a given setting of the fans and pump, the overall thermal resistance is relatively insensitive to the applied heat flux. Increasing the liquid flow rate improved the heat transfer to a point, after which there was diminishing returns in terms of higher power consumption. Lastly, increasing the fan speed had minimal impact on the overall heat transfer while significantly increasing power consumption and decreasing the COP. These findings highlight the importance of effective management of the fluid movers to optimize energy dynamics in liquid cooling systems. Optimising power consumption of the HLA system is crucial, as it remains the main drawback compared to conventional FFAC systems, highlighting the need for a comprehensive power optimization study.

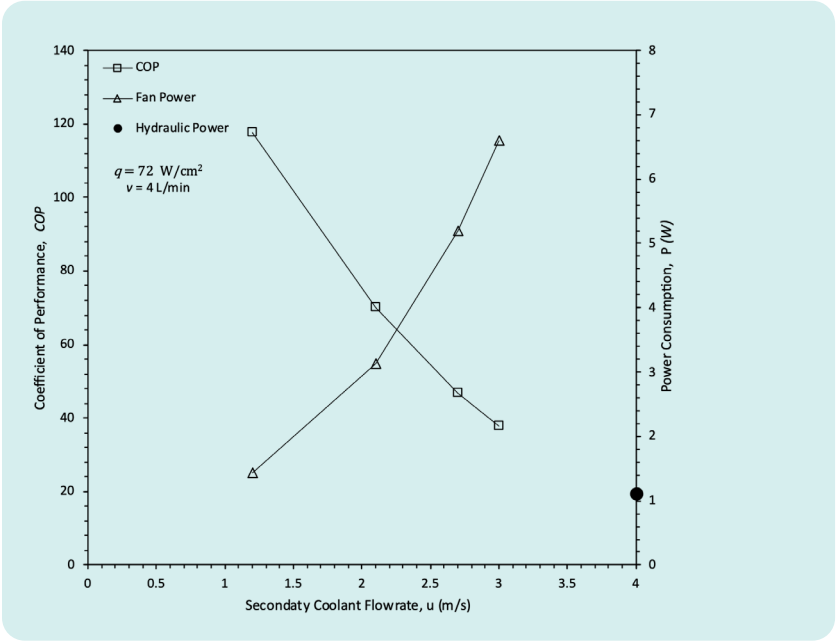

Fig 18. Power consumption of the liquid cooling system for varying air velocities measured at constant heat flux and liquid flow rate of 4 L/min: $T_{ambient}$ = 20°C.

3.4 Cooling Energy Optimization

To illustrate the potential for optimizing the cooling energy requirement of the HLA system, cases are considered where the fluid movers are controlled in such a way to maintain a constant die temperature. Practically, CPUs operate from idle to their TDP, and maintaining a constant die temperature would be beneficial since it would reduce thermal cycling over large computational load swings. Two case studies are considered, one for maintaining T_{die}=75°C and a more aggressive case where T_{die}=55°C, each for the more common situation where $T_{ambient}$=20°C. The full heat flux range from 6 W/cm^2 to 72 W/cm^2 is then considered for the primary liquid coolant flow rates between 0.5 L/min-8 L/min and the secondary air coolant velocities of 1 m/s -3 m/s.

Figure 19 (a) & (b), plot the total cooling power consumption to maintain die temperatures of 75°C and 55°C respectively, over the range of die heat fluxes and as a function of the pump and total fan flow required to sustain the setpoint temperature. To effectively maintain the constant die temperature at a given heat flux, Figure 19 shows that the minimum cooling power consumption adapts to a higher liquid flow rate and lower air velocity (fan speed). For example, considering the 75°C / 60W/cm^2 case, the same overall cooling performance can be achieved with a liquid flow rate of 0.52 L/min and air velocity of ~3.0 m/s, as it can with a higher liquid flow rate of 2.2 L/min and lower air velocity of 1.0 m/s. Here the total cooling power drops from 5.2 W to 1.8 W, which is better than what can be achieved by the FFAC system, despite the fact that the FFAC system cannot cool sufficiently at this heat flux level (Figure 6). Referring to the COP surface plot in Figure 20 (a), the resulting gain for this case is from 52.1 to 156.4, which is significant.

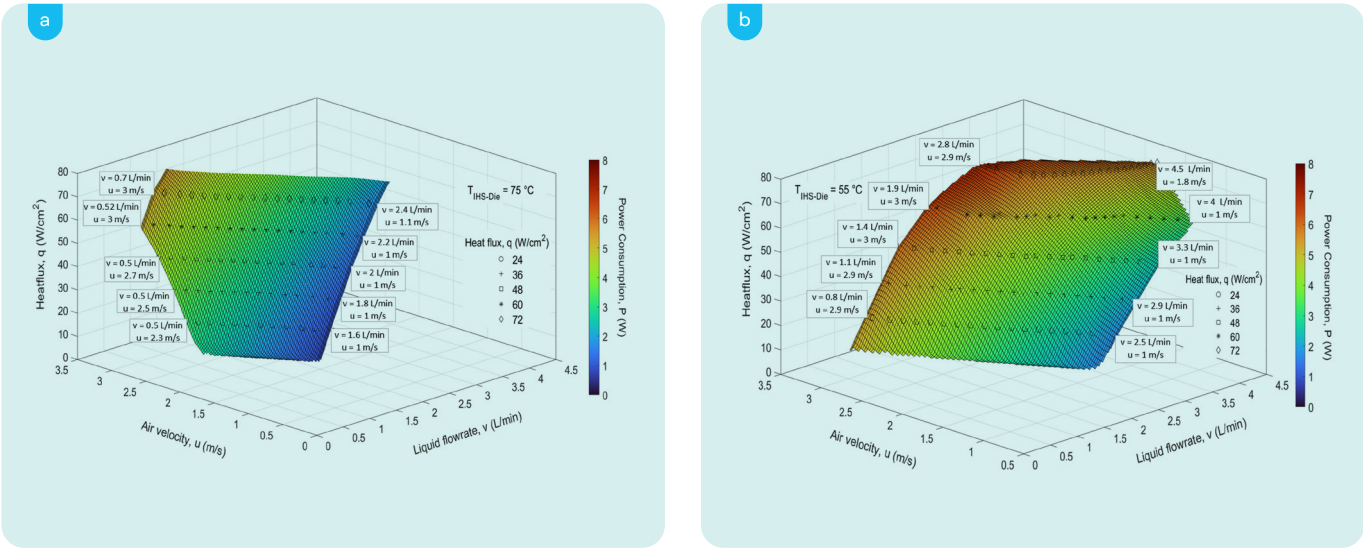

Fig 19. HLA power consumption surface plot for maintaining the die temperature at (a) T_{die} = 75 °C and (b) T_{die} = 55 °C, by varying liquid flow rate & air velocity with $T_{ambient}$ = 20 °C and R_{TIM1} = 0.1 K-cm2/W.

For the more aggressive cooling case of 55°C / 60W/cm^2, Figure 19 (b) and Figure 20 (b), it is optimal to operate the HLA system with the liquid and air fluid movers set at 4 L/min and 1 m/s respectively, resulting in an energy consumption of 3.3 W (COP~75). In contrast, utilizing the maximum air speed of 3 m/s, a liquid flow rate of only 1.9 L/min is required, though the energy consumption is 2 times higher, with a commensurate drop in COP. Therefore, opting for lower fan speed requires higher liquid flow rates to deliver equivalent cooling performance, though consuming comparatively less power, which improves the COP. It is interesting to note that the maximum air-speed of 3 m/s is never required for the two setpoint test cases, and apart from the highest heat flux, only ~ 1 m/s (lowest fan speed setting) is generally required for minimum cooling power and maximum COP. Thus, the temperature can be regulated, at minimum cooling power expense, by controlling the liquid flow rate.

In doing so, high COPs can be achieved for the range of CPU heat fluxes between 24 W/cm^2 to 72 W/cm2, ranging between COP~97 – 150 for the T_{die}=75°C setpoint, and COP~60 – 73 for the T_{die}=55°C setpoint. Importantly, this is comparable to a FFAC cooling system, whilst significantly expanding the maximum TDP capability of CPUs.

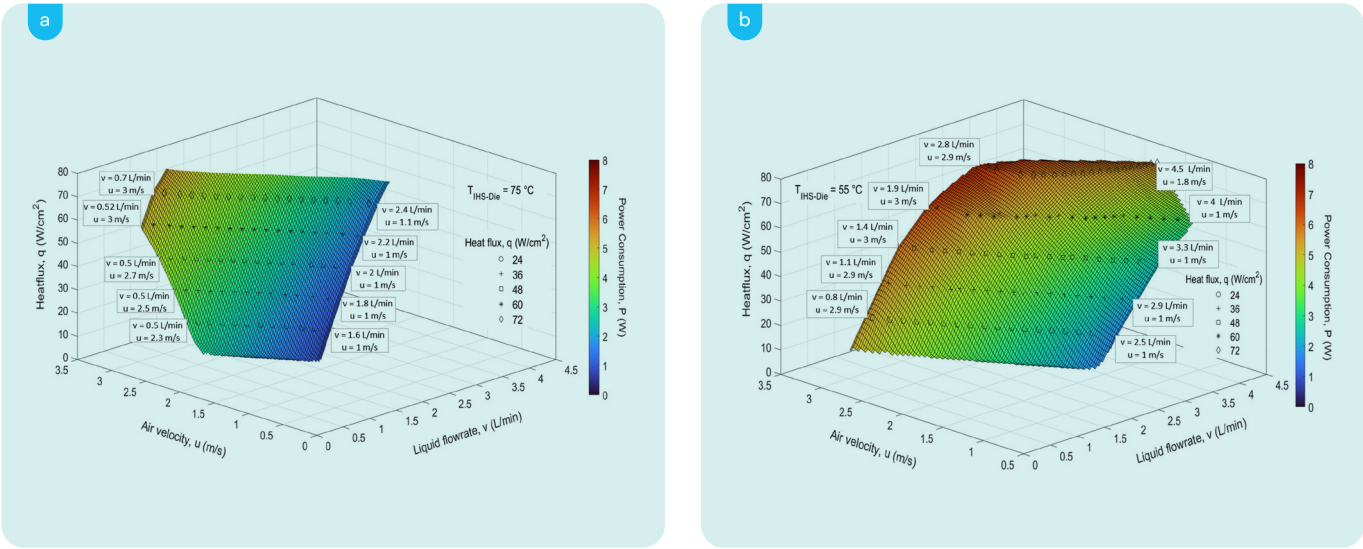

Fig 20. HLA COP surface plot for maintaining the die temperature at (a) T_{die} = 75 °C and (b) T_{die} = 55 °C, by varying liquid flow rate & air velocity with $T_{ambient}$ = 20 °C and R_{TIM1} = 0.1 K-cm2/W.

4. Conclusions

This case study aimed to investigate the cooling performance of a closed loop liquid cooled system for advanced thermal management of high-performance CPUs. For an emulated 2 x 2 cm^2 CPU die and range of input power and heat fluxes up to 288 W and 72 W/cm^2 respectively, the experimental investigation measured key cooling performance metrics of a Hybrid Liquid Air (HLA) system, such as die temperature and thermal resistance, over ranges of liquid flow rates (0.5 L/min – 8 L/min) and fan speeds (air velocities of 1 m/s to 3 m/s). The cooling pressure drop, and power penalty was also measured. To provide a comparison, a traditional Fan-Fin Air-Cooled (FFAC) system was also tested, and the following conclusions can be drawn;

(1) The HLA system outperforms the FFAC system in terms of source-to-sink thermal resistance and associated lower CPU temperatures.

(2) The feasibility limit of the FFAC system tested was around 42 W/cm^2 (168 W), whilst the HLA system has an expected capacity to cool over 112.5 W/cm^2 (450 W), illustrating that liquid-based cooling can support the escalating CPU power levels, whilst traditional air cooling cannot.

(3) The power requirement of the two fluid movers (pump and fans) in the HLA system can greatly exceed the FFAC system, if not correctly managed.

(4) With two fluid movers, maintaining a die temperature can be achieved by controlling either or both the pump and fan speeds, and the power cost penalty and overall COP of the system are sensitive to this choice.

(5) By controlling the pump flow rate and fan speed, a constant setpoint die temperature can be maintained over a range of CPU heat fluxes. Importantly, this can be achieved in such a way as to minimize the cooling power and thus maximize the system COP.

(6) For a given CPU heat flux, lowering the fan speed and increasing the liquid flow rate to achieve a die temperature setpoint minimizes cooling power and maximizes COP.

(7) For the most aggressive case investigated, with a CPU heat flux of 72 W/cm^2 (288 W) and CPU die temperature of T_{die}=55°C, the optimum power for the range of fluid flow rates tested is only 5.4 W, with an impressive COP=56, illustrating that industry concerns about the additional energy cost of liquid cooling can be largely addressed by engineering suitable control systems to manage the fluid movers.

The demonstrated superiority of the HLA system over the conventional FFAC system for the tested CPU heat transfer rates underscores the importance of advanced liquid cooling solutions for energy efficient and reliable operation of modern CPUs and other processors, such as GPUs. Future research will investigate means to further increase the cooling capacity of liquid-based cooling of high-powered processors by addressing the liquid coolant temperature in an HLA system. Being limited by the ambient air temperature, including a miniature vapour-compression refrigeration system has the potential to provide sub-ambient liquid coolant temperatures, thus increasing the cooling capacity of liquid-based thermal management systems.

Acknowledgements

This work was developed with financial support of the IRC-EPS Award EPSPG/2022/347 and CONNECT research centre via Science Foundation Ireland (SFI) grant number 13/RC/2077_P2.

References

[1] J.W. Elliott, M.T. Lebon, A.J. Robinson, Optimising integrated heat spreaders with distributed heat transfer coefficients: A case study for CPU cooling, Case Studies in Thermal Engineering. 38 (2022) 102354. https://doi.org/10.1016/j.csite.2022.102354.

[2] Intel product specifications, (n.d.). https://www.intel.com/content/www/us/en/ark.html (accessed May 16, 2023).

[3] M. Saini, R.L. Webb, Validation of models for air cooled plane fin heat sinks used in computer cooling, in: ITherm 2002. Eighth Intersociety Conference on Thermal and Thermomechanical Phenomena in Electronic Systems (Cat. No.02CH37258), 2002: pp. 243–250. https://doi.org/10.1109/ITHERM.2002.1012464.

[4] M. Saini, R.L. Webb, Heat rejection limits of air cooled plane fin heat sinks for computer cooling, IEEE Transactions on Components and Packaging Technologies. 26 (2003) 71–79. https://doi.org/10.1109/TCAPT.2003.811465.

[5] A.L. Moore, L. Shi, Emerging challenges and materials for thermal management of electronics, Materials Today. 17 (2014) 163–174. https://doi.org/10.1016/j.mattod.2014.04.003.

[6] D.B. Tuckerman, R.F.W. Pease, High-performance heat sinking for VLSI, IEEE Electron Device Letters. 2 (1981) 126–129. https://doi.org/10.1109/EDL.1981.25367.

[7] A. Mohammed Adham, N. Mohd-Ghazali, R. Ahmad, Thermal and hydrodynamic analysis of microchannel heat sinks: A review, Renewable and Sustainable Energy Reviews. 21 (2013) 614–622. https://doi.org/10.1016/j.rser.2013.01.022.

[8] A.A. Alfaryjat, H.A. Mohammed, N.M. Adam, M.K.A. Ariffin, M.I. Najafabadi, Influence of geometrical parameters of hexagonal, circular, and rhombus microchannel heat sinks on the thermohydraulic characteristics, International Communications in Heat and Mass Transfer. 52 (2014) 121–131. https://doi.org/10.1016/j.icheatmasstransfer.2014.01.015.

[9] Y.J. Lee, P.K. Singh, P.S. Lee, Fluid flow and heat transfer investigations on enhanced microchannel heat sink using oblique fins with parametric study, International Journal of Heat and Mass Transfer. 81 (2015) 325–336. https://doi.org/10.1016/j.ijheatmasstransfer.2014.10.018.

[10] G. Xie, J. Liu, Y. Liu, B. Sunden, W. Zhang, Comparative Study of Thermal Performance of Longitudinal and Transversal-Wavy Microchannel Heat Sinks for Electronic Cooling, Journal of Electronic Packaging. 135 (2013). https://doi.org/10.1115/1.4023530.

[11] Y.L. Zhai, G.D. Xia, X.F. Liu, Y.F. Li, Heat transfer in the microchannels with fan-shaped reentrant cavities and different ribs based on field synergy principle and entropy generation analysis, International Journal of Heat and Mass Transfer. 68 (2014) 224–233. https://doi.org/10.1016/j.ijheatmasstransfer.2013.08.086.

[12] E.G. Colgan, B. Furman, A. Gaynes, W. Graham, N. LaBianca, J.H. Magerlein, R.J. Polastre, M.B. Rothwell, R.J. Bezama, R. Choudhary, K. Marston, H. Toy, J. Wakil, J. Zitz, A practical implementation of silicon microchannel coolers for high power chips, in: Semiconductor Thermal Measurement and Management IEEE Twenty First Annual IEEE Symposium, 2005., 2005: pp. 1–7. https://doi.org/10.1109/STHERM.2005.1412151.

[13] J. Zhou, X. Cao, N. Zhang, Y. Yuan, X. Zhao, D. Hardy, Micro-Channel Heat Sink: A Review, J. Therm. Sci. 29 (2020) 1431–1462. https://doi.org/10.1007/s11630-020-1334-y.

[14] W. Qu, I. Mudawar, Experimental and numerical study of pressure drop and heat transfer in a single-phase micro-channel heat sink, International Journal of Heat and Mass Transfer. 45 (2002) 2549–2565. https://doi.org/10.1016/S0017-9310(01)00337-4.

[15] A.M. Kiper, Impinging water jet cooling of VLSI circuits, International Communications in Heat and Mass Transfer. 11 (1984) 517–526. https://doi.org/10.1016/0735-1933(84)90003-4.

[16] G.J. Michna, E.A. Browne, Y. Peles, M.K. Jensen, The effect of area ratio on microjet array heat transfer, International Journal of Heat and Mass Transfer. 54 (2011) 1782–1790. https://doi.org/10.1016/j.ijheatmasstransfer.2010.12.038.

[17] A.J. Robinson, E. Schnitzler, An experimental investigation of free and submerged miniature liquid jet array impingement heat transfer, Experimental Thermal and Fluid Science. 32 (2007) 1–13. https://doi.org/10.1016/j.expthermflusci.2006.12.006.

[18] A. Robinson, A Thermal–Hydraulic Comparison of Liquid Microchannel and Impinging Liquid Jet Array Heat Sinks for High-Power Electronics Cooling, Components and Packaging Technologies, IEEE Transactions On. 32 (2009) 347–357. https://doi.org/10.1109/TCAPT.2008.2010408.

[19] D. Nikolic, M. Hutchison, P.T. Sapin, A.J. Robinson, Hot spot targeting with a liquid impinging jet array waterblock, in: 2009 15th International Workshop on Thermal Investigations of ICs and Systems, 2009: pp. 168–173.

[20] A. Bhunia, C.L. Chen, On the Scalability of Liquid Microjet Array Impingement Cooling for Large Area Systems, Journal of Heat Transfer. 133 (2011). https://doi.org/10.1115/1.4003532.

[21] A.J. Robinson, R. Kempers, J. Colenbrander, N. Bushnell, R. Chen, A single phase hybrid micro heat sink using impinging micro-jet arrays and microchannels, Applied Thermal Engineering. 136 (2018) 408–418. https://doi.org/10.1016/j.applthermaleng.2018.02.058.

[22] J.S. Bintoro, A. Akbarzadeh, M. Mochizuki, A closed-loop electronics cooling by implementing single phase impinging jet and mini channels heat exchanger, Applied Thermal Engineering. 25 (2005) 2740–2753. https://doi.org/10.1016/j.applthermaleng.2005.01.018.

References

[23] S. Liu, J. Yang, Z. Gan, X. Luo, Structural optimization of a microjet based cooling system for high power LEDs, International Journal of Thermal Sciences. 47 (2008) 1086–1095. https://doi.org/10.1016/j.ijthermalsci.2007.09.005.

[24] J.-Y. Chang, H.S. Park, J.I. Jo, S. Julia, A System Design of Liquid Cooling Computer Based on the Micro Cooling Technology, in: Thermal and Thermomechanical Proceedings 10th Intersociety Conference on Phenomena in Electronics Systems, 2006. ITHERM 2006., 2006: pp. 157–160. https://doi.org/10.1109/ITHERM.2006.1645337.

[25] B.P. Whelan, R. Kempers, A.J. Robinson, A liquid-based system for CPU cooling implementing a jet array impingement waterblock and a tube array remote heat exchanger, Applied Thermal Engineering. 39 (2012) 86–94. https://doi.org/10.1016/j.applthermaleng.2012.01.013.

[26] Freecor FTC | Arteco Coolants, (n.d.). https://www.arteco-coolants.com/en/products/engine-coolants/x-oat-coolants/freecor-ftc (accessed March 31, 2023).

[27] R.L. Webb, Heat Exchanger Design Methodology for Electronic Heat Sinks, Journal of Heat Transfer. 129 (2006) 899–901. https://doi.org/10.1115/1.2717249.

[28] J.R. Taylor, An introduction to error analysis, University Science Books, Sausalito, Calif., 1997.

[29] R. Kempers, P. Kolodner, A. Lyons, A.J. Robinson, A high-precision apparatus for the characterization of thermal interface materials, Review of Scientific Instruments. 80 (2009) 095111. https://doi.org/10.1063/1.3193715.

[30] R. Kempers, A. Robinson, Heated Meter Bar Techniques: What You Should Know and Why, in: The Art of Measuring in the Thermal Sciences, CRC Press, 2020.

[31] Thermal Management for Intel® Xeon® Processors, Intel. (n.d.). https://www.intel.com/content/www/us/en/support/articles/000006710/processors/intel-xeon-processors.html (accessed January 20, 2023).

[32] How to Check and Monitor Your CPU Temperature, How to Check and Monitor Your CPU Temperature. (n.d.). https://www.avast.com/c-how-to-check-cpu-temperature (accessed July 2, 2023).

[33] Information about Temperature for Intel® Processors, Intel. (n.d.). https://www.intel.com/content/www/us/en/support/articles/000005597/processors.html (accessed May 30, 2023).

www.nexalus.com

CONTACT US:

Nexalus Ltd,
Unit 13, South Bank,
Crosse's Green,
Cork, Ireland,
T12 XT71

Nexalus Labs,
Water St,
Loughanalla,
Castlepollard,
Co. Westmeath, Ireland,
N91 EP90

Nexalus Research,
Trinity Research & Innovation,
O'Reilly Institute, Trinity College,
Dublin 2, Ireland,

info@nexalus.com
www.nexalus.com

Registered in Ireland No. 628880,
VAT No. IE3573607SH

WITH ENDORSEMENTS AND SUPPORT FROM:

ENTERPRISE IRELAND where innovation means business
Science Foundation Ireland sfi For what's next
Trinity College Dublin Coláiste na Tríonóide, Baile Átha Cliath The University of Dublin
CONNECT
IRISH RESEARCH COUNCIL An Chomhairle um Thaighde in Éirinn

QUALITY CERTIFICATIONS:

IATF 16949 / Automotive Standard
ISO 13485 / Medical Device Standard
ISO 50001 / Energy Management Standard
ISO 45001 / Occupational Heath & Safety Standard
ISO 14001 / Environmental Standard

NEXALUS
COOLING THE CLOUD