2025 24th IEEE Intersociety Conference on Thermal and Thermomechanical Phenomena in Electronic Systems (ITherm) | 979-8-3315-2428-9/25/$31.00 ©2025 IEEE | DOI: 10.1109/ITherm55376.2025.11235663

Revolutionary Thermal Solutions for Hot Chips

Ron Zhang, Laura Mirkarimi, Belgacem Haba, Gill Fountain, KM Bang, Suhail Sadiq and Arianna Avellan
Adeia, 3025 Orchard Parkway, San Jose, Ca 94513
Ron.Zhang@adeia.com

***Abstract* -As feature sizes shrink, power density increases as more functionality is added to semiconductor chips. Consequently, there is an increasing need for efficient heat extraction since inadequate cooling of semiconductor devices can lead to decreased performance, increased errors, and decreased lifespan. Providing adequate cooling can be especially challenging for semiconductor devices made using advanced process nodes, where features can be densely packed together, resulting in large heat generation in a relatively small area. Liquid cooling has long been one option for cooling semiconductor chips. However, there are limitations with conventional technologies.**

Today, the industry is considering immersion cooling in addition to Direct Liquid Cooling (DLC) with a cold plate. The former promises improved cooling performance compared with the latter but is challenging to implement at a large scale within existing infrastructure. This paper focuses on DLC and alternative approaches to managing higher power densities in hot chips. Conventional DLC solutions use a thermal interface material (TIM) to attach a cold plate to the IC die. Although the TIM improves interfacial heat conduction and mitigates CTE mismatch, it adds substantial thermal resistance along the heat conduction path. Some DLC solutions involve embedded microchannels in active devices, which offers significant thermal performance improvement and eliminates CTE mismatch between the cooling solution and the devices. However, several challenges exist that prevent widespread adoption of embedded microchannels, including high pressure drop, coolant leakage from the cooled devices, and the need to co-design the cooling channels and the IC device.

We have reconsidered the problem and have demonstrated a new integrated cooling solution (ICS) that addresses the challenges outlined above. Experimental test results show greatly improved thermal performance (~70% total thermal resistance reduction from junction to inlet with a power density in the range of 1.5W/mm² ~ 2W/ mm² for over 1kW die power) and significantly reduced pressure drop (like conventional copper cold plates). Unlike microchannel formation in the backside of an IC, this thermal solution is decoupled from the IC, eliminating any concern of coolant leakage. We share design optimizations through validated computational fluid dynamics (CFD) simulation, the performance data (thermal resistance & pressure drop), strategies for hot spot management, etc.

Keywords— liquid cooling, microchannel, thermal resistance, pressure drop

I. INTRODUCTION

The semiconductor industry is facing a revolution in compute power density increase, driven by the insatiable demand of AI architecture. As the power densities in processor integrated circuits increase, the thermal challenges for high performance computing in data centers is gaining attention among the industry. Typically, heat is removed from the IC by attaching a heat sink to the IC. The performance of a liquid based heat sink is defined as the thermal resistance, Θ_{JI}

$$\Theta_{JI} = \frac{\left(T_{\max,die} - T_{inlet\ liquid}\right)}{Power}$$

.

Thermal resistance consists of two main contributions, conduction and convection.

$$\Theta_{JI} = \Theta_{Conduction} + \Theta_{Convection}$$

The conductive resistance is virtually constant and is dominated by the TIM resistance. The convective resistance is inversely proportional to a function of the flow rate, Q; where α is a coefficient.

$$\Theta = \Theta_{conduction} + \Theta_{convection} = C + \alpha/f(Q)$$

The industry is moving from traditional air cooled solutions to direct liquid cooling, including immersion cooling. Currently, the industry standard in direct liquid cooling (DLC) involves a metal cold plate that is attached to the hot chips with a thermal interface material (TIM). The TIM is the largest contributor to the thermal resistance of heat conduction from the chip to the cold plate. For traditional metal cold plates, at high flow rates, the TIM resistance dominates as Θ_{conv} becomes relatively small.

We focus on reducing the thermal resistance of the system with our integrated cooling solution so that we can capitalize on that convective portion of cooling and efficiently remove heat from the hot chips to the cold plate where fluid may conduct the heat away from the chips. We use a silicon cold plate that is directly bonded to the hot chips, which spreads heat and reduces the thermal resistance. We effectively remove the conduction resistance by eliminating the TIM. Using a direct bonding technology, the IC is bonded to the cold plate at the molecular level forming Si-O-O-Si and/or Si-N-N-Si bonds and mixtures thereof at the interface [2]. We have evaluated a variety of cold plate designs to find an optimal thermal one defined by the lowest thermal resistance, given the same cooling area and flow rate. Initial efforts focused on thermal performance simulations. Next, prototypes were built for thermal testing to verify the predictability of the model with experimental results.

A. *Background in Si Microchannel Technology*

In 1980, Tuckerman and Pease reported that 50 micron wide microchannels operating in a laminar-flow regime, integrated within the silicon chip, showed a 40-fold improvement in thermal conductance, compared to conventional heat sink devices. The microchannel configuration was a face up configuration with the channels etched from the top surface (i.e., the backside of the die) leaving an unetched portion toward the integrated circuit. While these results were encouraging, the low power CMOS circuit development allowed lower power density chips to be manufactured, thus delaying the necessity to implement new thermal concepts. Moreover, the microchannels discussed in early papers are not easily fabricated and the channel geometry induced extremely high pressure drop, which are challenging for system integration. Various mitigation strategies were developed over the years to facilitate fabrication and reduce pressure drop [3-7]; however, no clear winner has emerged as a fully integrated solution which is compatible with the infrastructure used today in data centers.

In 2021, we began evaluating the microchannel technology to understand the thermal performance benefits and integration challenges for widespread adoption. Bringing our experience in direct and hybrid bonding technology to enhance advanced packaging performance may enable some new cooling approaches to emerge.

II. Design Analysis Through Simulation

3D models of the cold plate design were created in Ansys and subsequently solved with the IcePak module. Care was taken to ensure both the cold plate and the active die (represented by heat sources) were adequately represented in the model. A manifold was constructed to deliver flow to the cold plate. Material properties for the coolant (DI water in this case) and silicon were taken either from the built-in library or other verified sources. Boundary conditions were assigned to closely mimic the test setup. We used wall conditions for all of the surfaces on the boundaries of the fluid domain in the CFD model and free convection was applied to all other exposed surfaces. For the purposes of comparison, a chip size of 30mm x 30mm and 30mm x 26mm with a uniform power density were used for these simulations. For the 30mm x 30mm design, power density uniformity was maintained across multiple heating elements while for the 30mm x 26mm one, uniformity was maintained over the entire surface of the die.

A. *Cold Plate Design and Performance*

Several initial designs were evaluated with chip sizes of 20 mm x 20 mm to understand the tradeoffs in thermal performance between a cold plate with posts, elongated posts (rectangles) in both a straight and staggered pattern with respect to the direction of flow (Figure 1). In consideration of close proximity to the heat sources (active circuity), the ease of integration with a hot chip and overall mechanical reliability of the design, we selected a face down design with non-channel areas which mate to the manifold. Previous research groups such as Tuckerman and Pease have created microchannels or micro-texturing with large open cavity above these channels [1,3,4] which implemented a face-up microchannel design.

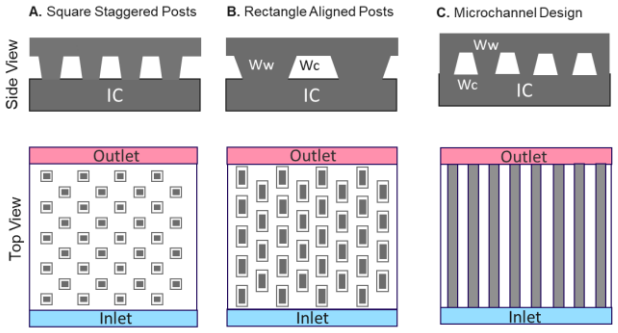

Fig 1: Schematic of various cold plate designs, including a) square staggered post, b) rectangle aligned post and c) micro channel.

The square post design had dimensions of 0.7mm x 0.7mm while the rectangular shapes had dimensions of 0.7mm x 2.5mm. Comparisons were made for the square posts to rectangle design to understand the influence of more Silicon interaction with the cooling fluid. The aligned rectangular post design creates a straight channel for flow and a pressure drop that was 4x better than the staggered layout. However, the maximum temperature was reduced by 4 degrees in the staggered design (Table 1). To further explore the reduction of temperature with more fluid interaction with the cold plate, we examined a two layer silicon cold plate structure where the liquid moves over the rectangular posts in both horizontal and vertical direction as shown in Figure 2. The maximum temperature was significantly reduced to 64C however, the pressure drop increased by 750x, making this concept less preferable for industry adoption.

TABLE 1: MAXIMUM DIE TEMPERATURE FOR VARIOUS COLD PLATE DESIGNS

Design	Type	Inlet Velocity (m/s)	Max Die Temperature (oC)	Pressure Drop (PSI)
Rectangle	Aligned	1.5	105	0.02
Posts	Staggered	1.5	101	0.09
Vertical	Staggered	1.5	64	15
Full Channel	Aligned	1.5	101	0.01

Next we focused our attention on a microchannel design that spans the length of the chip. In the initial microchannel design, we selected a channel size of 0.7mm x 0.5mm depth with one inlet and one outlet at the two edges as shown in the

schematic in Figure 1c. The microchannel design showed a similar maximum temperature to the staggered post design with a 9 fold improvement in pressure drop. Based on these results, we continued to focus on the microchannel design optimization (1c).

Examination of the temperature contour in a given cross section of a trapezoidal channel reveals that only a small

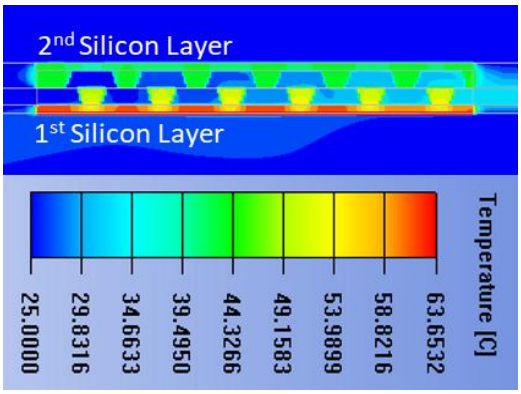

Fig 2: Two-layer Si cold plate design noted as vertical staggered type in Table 1.

region within close proximity to the liquid to silicon interface is heated while the rest remains unheated with temperatures close to the inlet liquid, indicating little heat transfer from the interface to the interior of the liquid. To enlarge the heat transfer region relative to the cross sectional area, we moved from a trapezoidal channel to a triangular channel.

While comparison among heat sink designs is challenging due to system level differences, it is important to assess performance differentiation with design concepts of the cold plate. To achieve this, we scaled our design to the chip size of the original Tuckerman and Pease paper and maintained the same nominal average channel velocity of the fluid flow by keeping the same cross sectional area and preserving the unique geometrical feature of the two designs. The thermal resistance of our optimal design is plotted as a function of the flow rate (Figure 3). The enhanced design in our cold plate shows approximately 14% reduction in thermal resistance over the original concept at lower flow rates. At higher flow rates, and high pressure drop, Tuckerman and Pease geometry looks better. These results emphasize the importance of optimizing the microchannel design for thermal performance in a particular application.

Unlike Tuckerman and Pease who used a "face up" microchannel design, our design is "face down" with the liquid in direct contact with the backside of an active die. This feature moves the coolant as close as possible to the active circuity without physically modifying the active die. Additionally, our integrated cooling solutions is easier for

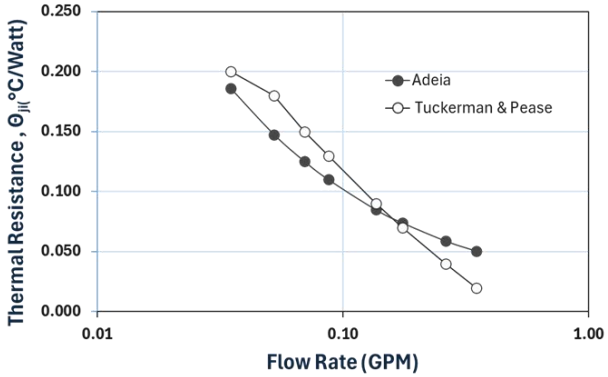

Fig 3: Adeia cold plate design compared to Tuckerman and Pease.

integration with integrated circuits and multi-die modules within the industry due to the fabrication process and compatibility with the Si supply chain.

III. Experimental Verification

In order to validate and calibrate our models, we built prototypes to test the performance of the various Si microchannel design and compared those to several off-the-shelf metal based cold plates with commonly available TIM materials.

A. Prototype Details

Thermal resistor chips were fabricated with standard BEOL silicon processes to create both 30mm x 30mm and 30mm x 26mm heater chips. The insulating dielectric covering the resistors was removed in the area where the heater was soldered to the power lines. The microchannel cold plates were fabricated with standard BEOL processes to create the microchannel features shared in Figure 1c. The integrated manifold was created with 3D printing and secured to the cold plate with an insulating adhesive. The inlet and outlet fixtures were ¼ inch compression fittings that allowed connectivity to the cooling fluid delivery system.

B. Test Apparatus and Experimental Details

The test apparatus includes a fluid delivery system, flow meters, pressure sensors, power supplies, thermocouples and IR cameras. The flow rate was controlled to within 0.01gallons per minute (GPM) by an ultrasonic flowmeter (DUK-11N4GC3T0). The water pressure and temperature were also measured at the inlet and outlet of the prototype using pressure meters and thermocouples. These readings were recorded for each test condition.

Temperature profiles of the prototype were recorded with an IR camera (FLIR E76) that captured the infrared radiation of the prototype and converted this radiation into an electrical signal. Then, the camera's processor converted this electrical signal into a temperature reading. The accuracy of the recordings was ± 2ºC. Before imaging the prototype, it was necessary to coat the prototype surface to ensure minimum reflectivity, to improve the accuracy of thermal imaging. The

prototype was coated to ensure the emissivity value is high, to maintain a low reflectivity from the surface. In our case, the emissivity was calibrated and set to 0.98, which correlates quite well with the thermocouple readings. Once this calibration was completed and the prototype was coated, it was ready for data collection. A 2000-Watt power supply was attached to the heater resistor chip. It was connected to the prototype through solder joints. Each resistor was connected to a channel on the power supply and controlled independently of the other resistors. A safety control box was installed in the form of a switchbox between the power supply and the prototype to switch the power on/off for each individual resistor. The power settings were held constant across all resistors during the test to achieve uniform power density.

The experimental protocol to create the power vs maximum temperature measurements in the prototype was to adjust the power, wait for system equilibrium, collect the data and repeat. Typically, each prototype part was tested a minimum of three full power cycles so that each data point represents an average of those measurements. Flowrates and power measured varied between 0.2-1.5 GPM and between 250 – 2000W, respectively. In the camera software, the area of the prototype was selected and the highest and lowest temperature measured in this area was recorded for 35 seconds. This recording contained over 1000 data points which were then averaged. This was repeated for every test condition at least three times to ensure reproducibility and repeatability (gage R&R). With the power, inlet water temperature and maximum die temperature values, the thermal resistance of the prototype system, Θ_{JI}, was calculated using the equation in the Introduction section of this paper and plotted as a function of flowrate.

Two off-the-shelf metal cold plates were selected to compare performance to our ICS module. A 50mm, Lazmin CPU copper cooling block (purchased through Amazon, UPC 750086360863), was attached to the resistor heater chip with a silver-based TIM as commonly used for a gaming machine. Additionally, a high-performance compute cold plate was selected that fit the designed heater chip. A Honeywell PTM7950 TIM was also used to improve conductive thermal resistance between the chip and the cold plate.

C. Experimental Results

The thermal resistance of the prototype system was measured for a variety of prototype designs. For comparison, the thermal resistance of several integrated cooling solution (ICS) designs are compared to an example gaming cold plate and a representative high performance compute (HPC) metal cold plate all attached to the same heater chip and measured at the same fluid flow rate of 0.2 GPM (Figure 4). We arbitrarily selected a range of flow rates to keep the pressure drop below 4.5 psi. The high performance cold plate showed a reduction in thermal resistance of about 10% compared to the gaming cold plate. Our ICS-TV5 design showed an improvement over the gaming cold plate of approximately 66%. Further improvement in the channel design allowed us to reduce the thermal resistance to approximately 79%. The thermal resistance of the ICS cold plates continue to improve with increased flow rate and remain significantly lower than the metal based cold plate comparisons.

The impact of lowering the thermal resistance of the system is that a much lower maximum temperature can be achieved with the same cooling flow rate. In Figure 5, the maximum temperature for gaming, HPC and ICS cold plate are compared for the same thermal power delivered across the chip. In this graph, we see that the 70% thermal resistance reduction in the ICS cold plate translates to the ability to deliver ~ 3x power density to a chip to maintain a temperature of ~100C.

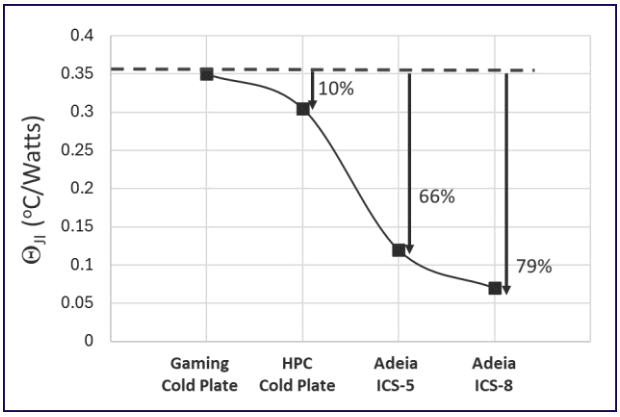

Fig 4: Measured thermal resistance is shown with a constant flow rate of 0.3 gallons per min. for a) gaming cold plate, b) high performance compute metal cold plate, c) ICS-5 0 and d) ICS-8 Si cold plate.

In Figure 5, with a cooling fluid flow rate of 0.3GPM, the gaming cold plate has a temperature of 102C at a power of 240 Watts, while the ICS has a temperature of 102 at 720Watts.

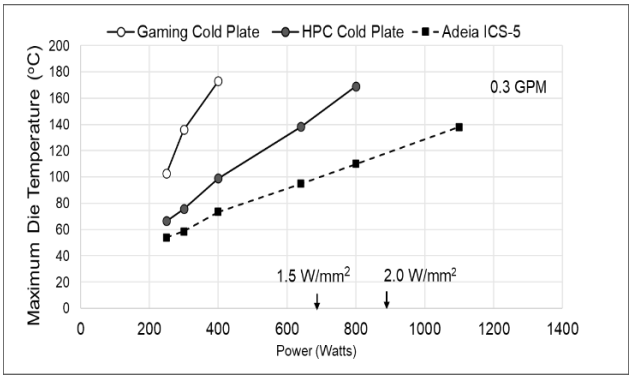

Fig 5: Measured maximum temperature as a function of power for a) gaming cold plate, b) high performance compute metal cold plate, and c) ICS-5 Silicon cold plate with 0.3GPM liquid flow rate.

In comparison, the high-performance cold plate closes the performance gap somewhat; however, the ICS cold plate, can sustain a power density increase of 1.75x compared to the HPC cold plate.

The pressure drop across prototype heat sinks was held below 4.5 psi at this low flow rate. Specifically, pressure drop for the gaming cold plate was 0.2 psi, the HPC cold plate was 1psi while the ICS-5 and ICS-8 were 4.3 psi and 1.6 psi, respectively.

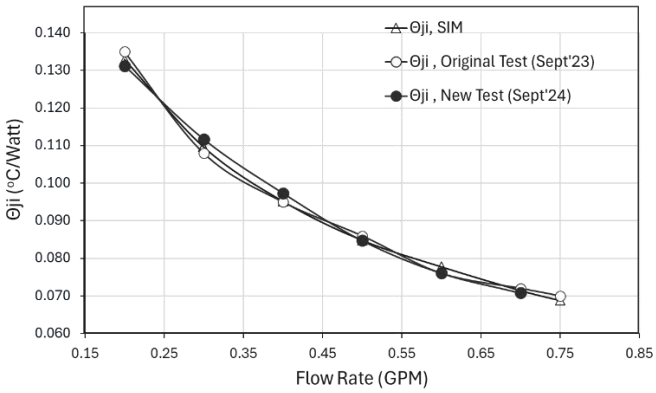

Fig 6: The simulation results of the thermal resistance are plotted as a function of cooling fluid flow rate for ICS-5 to experimental data taken a year apart.

This direct comparison above shows the enhanced performance of a silicon microchannel cold plate at relatively low flow rates to efficiently cool a hot chip. As Tuckerman and Pease explain in their seminal paper, the microchannel cold plate continues to improve the thermal resistance at higher flow rates, which is different from conventional cold plates with a TIM layer. The latter tend to reach a cooling limit much faster, i.e., further increase in the flow rate won't reduce thermal resistance.

To understand the predictive ability of our simulation protocol, we plot the simulated thermal resistance of ICS-TV5 with the experimental data taken in November 2023 and a year later in 2024 (Figure 6). The data indicates that our simulation protocol is predictive to within 5% of our experimental performance of the prototypes. Therefore, we have confidence that we may use CFD to perform cold plate design optimization and system considerations. Additionally, the results measured over a year separation for the prototype indicate the rigor in the experimental test protocol.

IV. DISCUSSION AND CONCLUSION

We have shared our "face down" cold plate designs and discussed the implications of design on thermal performance. Using CFD, we completed multiple design of experiments (DOEs) to identify improved silicon cold plate designs for our integrated cooling solution. The early simulation results showed that a full microchannel with a trapezoidal cross section had much lower pressure drop compared to a microtextured surface of a Si cold plate (TABLE 1). Results from both trapezoidal (ICS-5) and triangular cross section microchannel (ICS-8) shapes were shared with their respective thermal resistance reduction obtained with the CFD simulations (Fig. 4). The measured thermal resistance of the gaming and HPC cold plates do not allow us to easily separate out the contribution from the TIM or the heat sink without elaborate instrumentation and we didn't pursue it. However, the thermal interface resistance for our integrated cooling solution is negligible, thereby adding minimal resistance to the system allowing us to substantially reduce the total thermal resistance as seen in the enhanced performance.

Recently, a revival of microchannel concepts like the original proposals of Tuckerman and Pease have surfaced in a few papers; however, the specific dimensions of the cold plate or system are not readily available for comparison [8]. Therefore, we created a comparative model of the ICS-8 structure to the Tuckerman and Pease structure, for the purpose of evaluating the difference in the thermal resistance of the cold plate design while keeping all other system level constraints constant. The thermal resistance of our optimized cold plate design was found to have a 14% lower resistance compared to the original microchannel cold plate design (Fig 3)[1].

Experimental thermal testing of the ICS prototype designs verified the performance enhancement trends predicted in simulation to within 5% of measured data. Figure 6 shows the experimental and simulated data plotted on top of each other. A direct comparison of the ICS prototype heat sinks to representative gaming and HPC cold plates was completed to understand the relative performance. The advantages of the new integrated cooling solution for the same heater chip power densities and cooling liquid flow rates were due to much lower thermal resistance and the optimized design of the cold plate. The enhanced performance enables lower junction temperatures, higher device reliability and reduced cooling costs.

One important aspect that we have not discussed in detail here is pressure drop. As traditional microchannel cold plates have relatively small cross sections, they tend to induce a higher pressure drop through the channels. We tested all parts at low flow rates where the comparative pressure drop between the conventional cold plates and the ICS structures is very similar, ranging from 1psi to 4.3 psi. The design differences in the ICS structures to achieve 4.3 or 1.6 psi pressure drop are related to the presence of more inlet/outlet configurations. The pressure drop mitigation is influenced by the integrated manifold, cold plate design and inlet outlet configurations in the system, which we will cover in future dedicated articles.

Hot spots are inevitable in actual integrated circuits despite efforts to achieve uniform power density over an entire die. Additionally, the application of this technology is most likely in the HPC module market which has multiple die and or stacked die packages to consider. Future papers will share our mitigation strategies to achieve the optimal balance between thermal performance and pressure drop for multi-die system level configurations.

REFERENCES

[1] D.B.Tuckerman and R.F.W. Pease, IEEE Electronic Device Letters, EDL-2, No. 5, May 1981.

[2] G. Gao *et al*., "Development of Low Temperature Direct Bond Interconnect Technology for Die-To-Wafer and Die-To-Die Applications-Stacking, Yield Improvement, Reliability Assessment," *2018 International Wafer Level Packaging Conference (IWLPC)*, San Jose, CA, USA, 2018, pp. 1-7.

[3] E.G. Colgan, K. Schleupen, P. Mann, R.P. Kuder, J. Woong Nah, K. Sakuma, V. Alvarez, L. Turlapati, Danny C. Wong, Stanislav Lukashov, Steve Cordes, and James L. Speidell; "Fabrication and Performance of 300-mm Wafer-Scale Silicon Microchannel Cooler", IEEE Transactions on Comp, Pkg and Manf. Tech., Vol. 3 No. 14, (2023).

[4] G. M. Harpole and J. E. Eninger, "Micro-channel heat exchanger optimization," in *Proc. 7th IEEE Semiconductor Thermal Meas. Manage. Symp.*, Feb. 1991, pp. 59–63, doi:

[5] R. L. Webb, "Effect of manifold design on flow distribution in parallel micro-channels," in *Proc. ASME Int. Electron. Packag. Tech. Conf. Exhib., Int. Electron. Packag. Tech. Conf. Exhib.*, Maui, HI, USA, vol. 2, Jul. 2003, pp. 527–535.

[6] F. Laermer and A. Urban, "Challenges, developments and applications of silicon deep reactive ion etching," *Microelectron. Eng.*, vols. 67–68, pp. 349–355, Jun. 2003.

[7] F. Laermer, S. Franssila, L. Sainiemi, and K. Kolari, "Deep reactive ion etching," in *Micro and Nano Technologies, Handbook of Silicon Based MEMS Materials and Technologies*, M. Tilli, M. Paulasto-Krockel, M. Petzold, H. Theuss, T. Motooka, and V. Lindroos, Eds., 3rd ed. Amsterdam, The Netherlands: Elsevier, 2020, ch. 16, pp. 417–446.

[8] Y.J. Lien, C.C. Hsieh, T. Ku, L. Wang, P.J. Chen, K.C. Yee, C.H.D. Yu, "An Energy-efficient Si-integrated Micro-cooler for High Power and Power-density Computing Applications", IEEE 74th Electronic Components and Technology Conference (ECTC), (2024), p. 1025.