MULTIPLATFORM NANOFILM EVAPORATION HEAT EXCHANGER

A LEAP IN HIGH POWER IC THERMAL PERFORMANCE

CONTENTS:

EXECUTIVE SUMMARY:

As the semiconductor manufacturing process evolves, more transistors are integrated, thus pushing the limits of higher computation power. As the density of heat increases for each newer generation of processors, keeping them at a reasonable working temperature is crucial for higher frequencies, longevity, and stability. Especially for data centers and high-performance computing, existing cooling technologies are struggling to keep up with the increased heat density and TDP. Consequently, the transition from air to liquid cooling becomes viable as liquid is more efficient at heat removal. We believe liquid cooling is the most efficient and elegant cooling strategy for large-scale deployment. Among the barriers towards large-scale industrial adoption, cooling performance is the most significant but challenging factor. Therefore, we developed a nanofilm evaporation heat exchanger that to dissipate heat flux over 500W/cm^2. We use data from existing studies that tested the same processor model for comparison to demonstrate the performance of our solution. At the end of the white paper, we discuss the promising application of our solution in the field of data center, HPC and consumer electronics thermal management.

INTRODUCTION:

Moore's Law has long driven the advancement of the semiconductor industry, propelling it towards higher transistor density and improved performance. (figure 1) With the stagnation of Dennard scaling, increased transistor density also entails a higher density of heat. The demand for highly efficient heat dissipation to accommodate the increased power draw rises correspondingly. Additionally, as it becomes more difficult to shrink the size of transistors, the trend towards 2.5D and 3D packaging will become inevitable, further underscoring the importance of cooling performance for high-performance processors.[1] Conventional methods like air cooling heat pipes and microchannel cold plates are struggling to keep up with the need. Therefore, introducing better and more efficient ways of heat dissipation will be the key to unlocking higher computational power.

Evaporation at the solid-liquid-vapor interface has shown great potential for highly efficient heat removal. Despite being the subject of extensive academic research over the past decades, there still remains inadequate fundamental understanding of evaporation for applications in cooling devices. Consequently, existing phase-change-based products suffer from low efficiency and performance. We have researched and implemented our unprecedented nano-film evaporation cooling technology in a vapor chamber configuration for high-performance processors. We are demonstrating its performance through benchmarking and overclocking an Intel i7-8700K CPU.

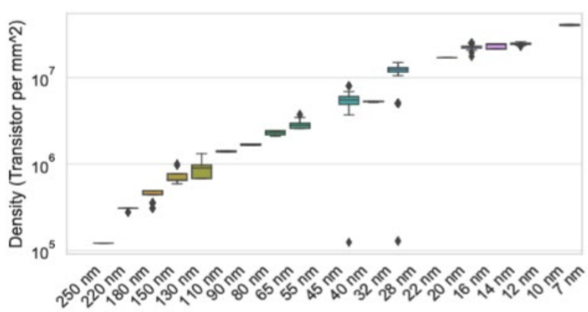

FIG 1: Transistor density increases for each generation [6]

TEST SETUP

In this work, we chose the Intel i7 8700K as the heat source rather than TTV. Although TTV offers more flexibility, for example, the heat input can be accurately controlled, it may not reflect a heat sink's real performance under deployment scenarios, as the distribution of cores and corresponding heat flux varies for different processors. Therefore, it is more realistic and representative to benchmark on real CPUs and compare their performance with other cooling methods.

From a thermal perspective, the Intel i7 8700K has a high heat flux, making it challenging to cool, more than 90% of power is concentrated on a 54 mm^2 core area. We list our calculated heat flux of the i7 8700K in figure 5. It's estimated to have over 300W/cm^2 heat flux when core power reaches 150W. To better understand the heatsink performance, we removed the IHS from the stock processor, as it usually helps in conducting heat. As shown in the figure 4 and figure 7, our heatsink has direct contact with the processor die, and the fins are water cooled to maintain a constant 23°C liquid temperature. A thin layer of liquid metal was used as the thermal interface material. We tested our testbench with Cinebench R20 and recorded data including frequency, core power, and temperature using HWINFO64. We used the Intel Extreme Tuning Utility to overclock the processor.

FIG 2: Intel 8700K before and after removing the IHS

	Intel i7 8700K
Socket	FCLGA1151
#Cores	6
Die Area	149 mm^2
Core Area	54 mm^2
Base Clock	3.7 GHz

Base TDP	95 W

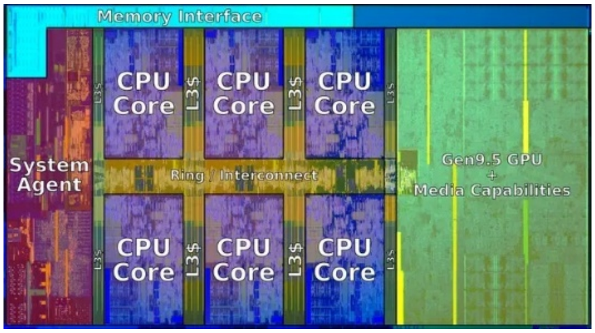

FIG 3: I7 8700k Die Shot

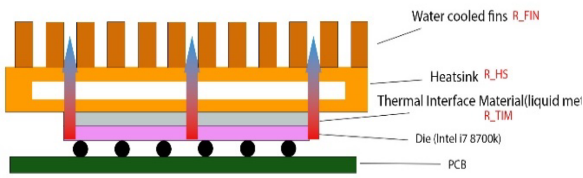

FIG 4: Heat transfer Diagram & Thermal Resistance

$$R_{TH}=R_{TIM}+R_{HS}+R_{FIN}$$

Heat In (Watt)	Density of Heat (W/cm^2)
100	207.1953
110	227.9148
120	248.6343
130	269.3538
140	290.0734
150	310.7929
160	331.5124
170	352.2319
180	372.9515
190	393.671
200	414.3905
210	435.11
220	455.8296

FIG 5: Estimated Density of Heat

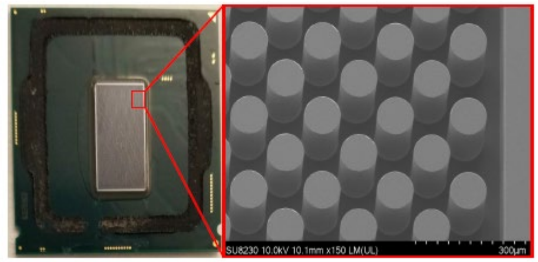

FIG 6: The integrated silicon microchannel solution [2]

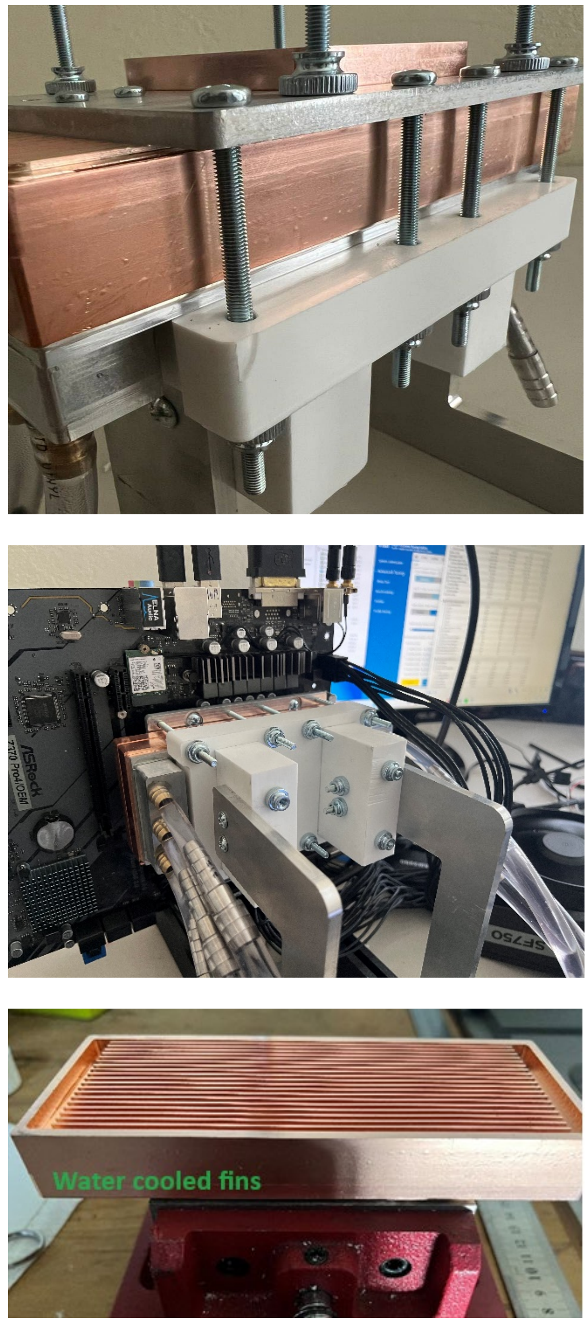

FIG 7: Test Rig Setup

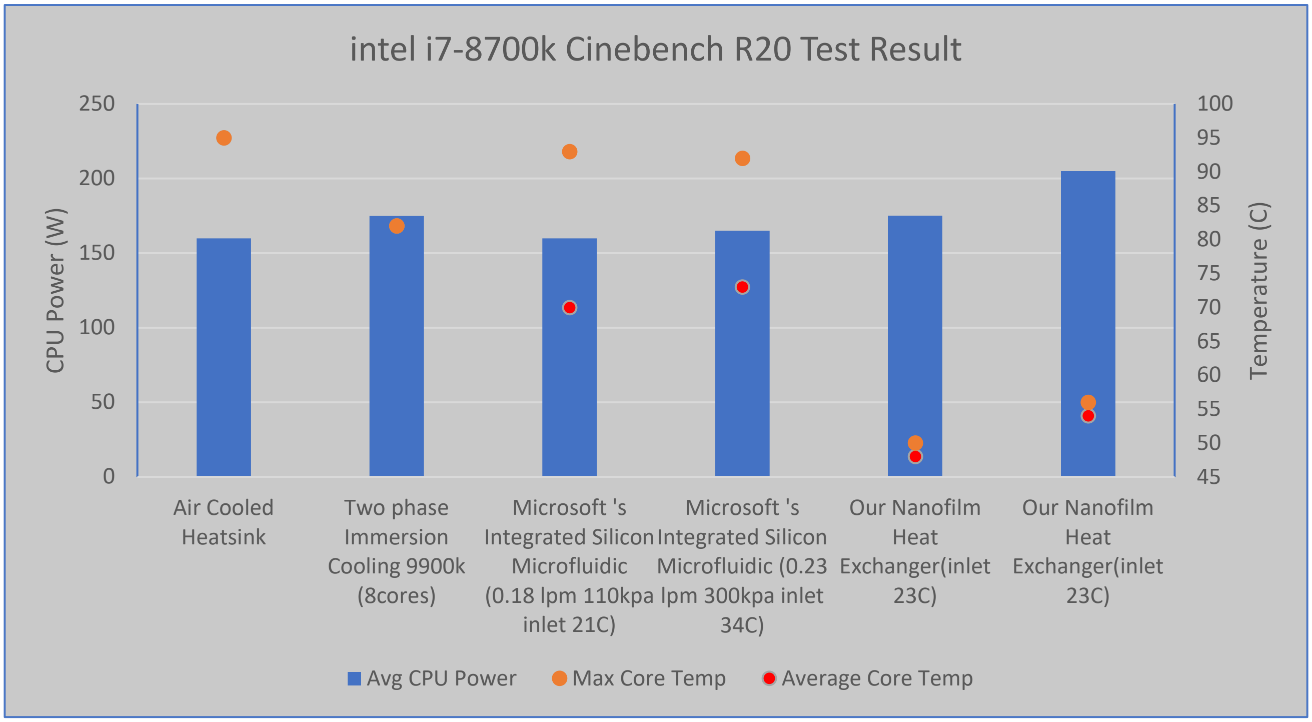

FIG 8: Comparison of multiple cooling techniques with i7-8700k [2]

RESULTS:

In figure 8, we show the core temperature of running the Cinebench R20 benchmark at different watts. Our nanofilm heat exchanger records a 48 °C average core temperature, 50 °C max core temperature at 170W CPU package power. In figure 8, the bottom-line air-cooler thermal throttled at 160W. We compare our result to the integrated silicon microfluidic heatsink developed by Microsoft & Georgia Institute of Technology in figure 6 [2], because it's considered an extremely efficient forced convection liquid cooling embodiment. It would reflect the optimum performance direct-to-chip liquid cooling could achieve. By comparison, our solution achieves 22 °C lower average core temperature under the same load at 170W. When the flow rate is tuned up and the pressure drop reaches 300Kpa, the integrated silicon microfluidic method has a boost in performance achieving 73 °C at 170W with 34C inlet temperature. The Open Compute Project (OCP) Cold Plate specification advises the node inlet pressure to be less than 300 Kpa. [4] So even it's possible to gain extra performance by increasing the water flow, the pressure drop would be too high for real world use cases.

Noting that the Integrated silicon microfluidic heatsink was made directly on the silicon die therefore eliminates the thermal resistance (R_{TIM}) of thermal interface material, which is not the case for our demo. The heat resistance of our case already takes the R_{TIM} and R_{FIN} into calculation, and our solution still has much lower thermal resistance compared to other techniques shown in figure 9.

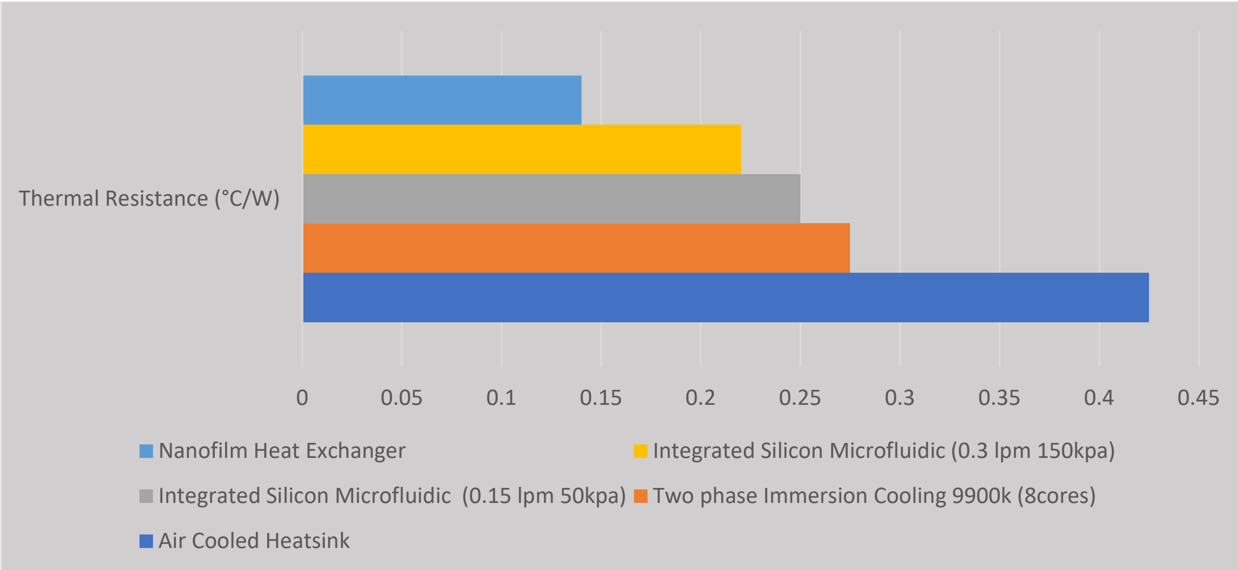

FIG 9: Thermal Resistance of different cooling techniques

APPLICATIONS:

Currently, with increasing power consumption and higher TDP, data centers and HPC are transitioning towards liquid cooling.[3] Cold plates and immersion cooling are the two main approaches. Our heatsink is an ideal choice for either single phase liquid cooling or single-phase immersion cooling, providing unprecedented thermal performance and unlocking new possibilities. Our nanofilm evaporation vapor chamber has a heat transfer coefficient (40 W/cm^2K) on the evaporator that is six times higher than conventional methods like heat pipes or microchannel cold plates (most less than 5 W/cm^2K). This means it requires a very small delta T to spread the concentrated heat from the heat source to a large condensation structure. Consequently, the attached fins can be easily cooled according to the user's preference.

Nowadays, the most popular implementation of server liquid-cooling uses microchannel cold plates directly attached to heat dissipating components like CPUs and GPUs with in the server to direct the flow of liquid coolant Due to the high efficiency of our technology, we can replace the microchannel cold plates with our vapor chamber and fins. This embodiment eliminates the high-pressure drops and clogging issues commonly found in microchannel cold plates, as the fins can be made much larger and cover a wider area. According to Meta, the worst-case scenarios of coolant pressure drop changes showed a 50% increase in pressure drop after 250 days of operation.[7] Since the microchannels are tiny, particles can easily accumulate and clog them. With larger and wider fins, this would not be a problem, and changing the working fluid from water to other types of dielectric working fluids becomes possible.

Immersion cooling has commonly been criticized for its unfavorable thermal performance. For example, *"single phase immersion using low vapor-pressure oils is also becoming increasingly common in data center cooling, but it does not really compete with single phase water liquid cooling in terms of heat transfer performance.*" [5]

However, with the technology we provide, this is no longer the case. Our heatsink offers superior thermal performance compared to cold plates and existing two-phase immersion cooling. A sample implementation of the heat exchanger for a 2U server, as seen in Figure 10, is a perfect fit for single-phase immersion and can be customized as needed.

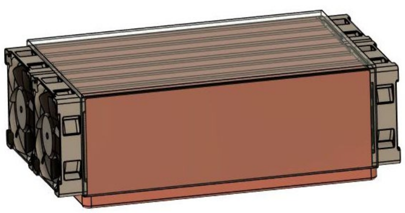

FIG 10: Purposed Heatsink with optional fans or water cooling

In addition, our technology can be adopted in consumer electronics thermal management and can be miniaturized for use in mobile devices such as laptops, tablets, and phones. For desktops, we are preparing to launch our first-generation CPU cooler by the end of Q4.

REQUEST FOR DEMO & COLABORATION:

In this white paper we demoed the benchmark of our solution with Intel i7-8700k. We understand different users may have different needs, and some readers may have concerns about our technology. So, we encourage our readers to contact us then we are more than pleased to make live demos specific to your interest, and we believe a working demo speaks for itself.

REFERENCES:

1. B. Ramakrishnan et al., "CPU Overclocking: A Performance Assessment of Air, Cold Plates, and Two-Phase Immersion Cooling," in IEEE Transactions on Components, Packaging and Manufacturing Technology, vol. 11, no. 10, pp. 1703-1715, Oct. 2021

2. S. Kochupurackal Rajan, B. Ramakrishnan, H. Alissa, W. Kim, C. Belady and M. S. Bakir, "Integrated Silicon Microfluidic Cooling of a High-Power Overclocked CPU for Efficient Thermal Management," in IEEE Access, vol. 10, pp. 59259-59269, 2022

3. M. Jalili et al., "Cost-Efficient Overclocking in Immersion-Cooled Datacenters," 2021 ACM/IEEE 48th Annual International Symposium on Computer Architecture (ISCA), Valencia, Spain, 2021, pp. 623-636

4. Open Compute Project, "WHITE PAPER: AN ADVANCED LIQUID COOLING RACK DESIGN FOR DATA CENTER"

5. Austin M. Shelnutt, "Two-Phase, or Not Two-Phase? Strengths and weaknesses of two-phase cooling in the Data Center", Strategic thermal labs, LLC

6. Sun, Yifan & Agostini, Nicolas & Dong, Shi & Kaeli, David. (2019). Summarizing CPU and GPU Design Trends with Product Data.

7. Yin Hang et al., "Long-Term Reliability Test on An Air Assisted Liquid Cooling System", May 2024, META

Contact: info@entroflowtech.com