WHITE PAPER

A COMPARATIVE STUDY AND IN-DEPTH ANALYSIS OF COOLING METHODS THROUGH OPTIMAL COLD PLATE DESIGN

Real-World Experimental Insights into Cooling 700-Watt GPUs

Rev A

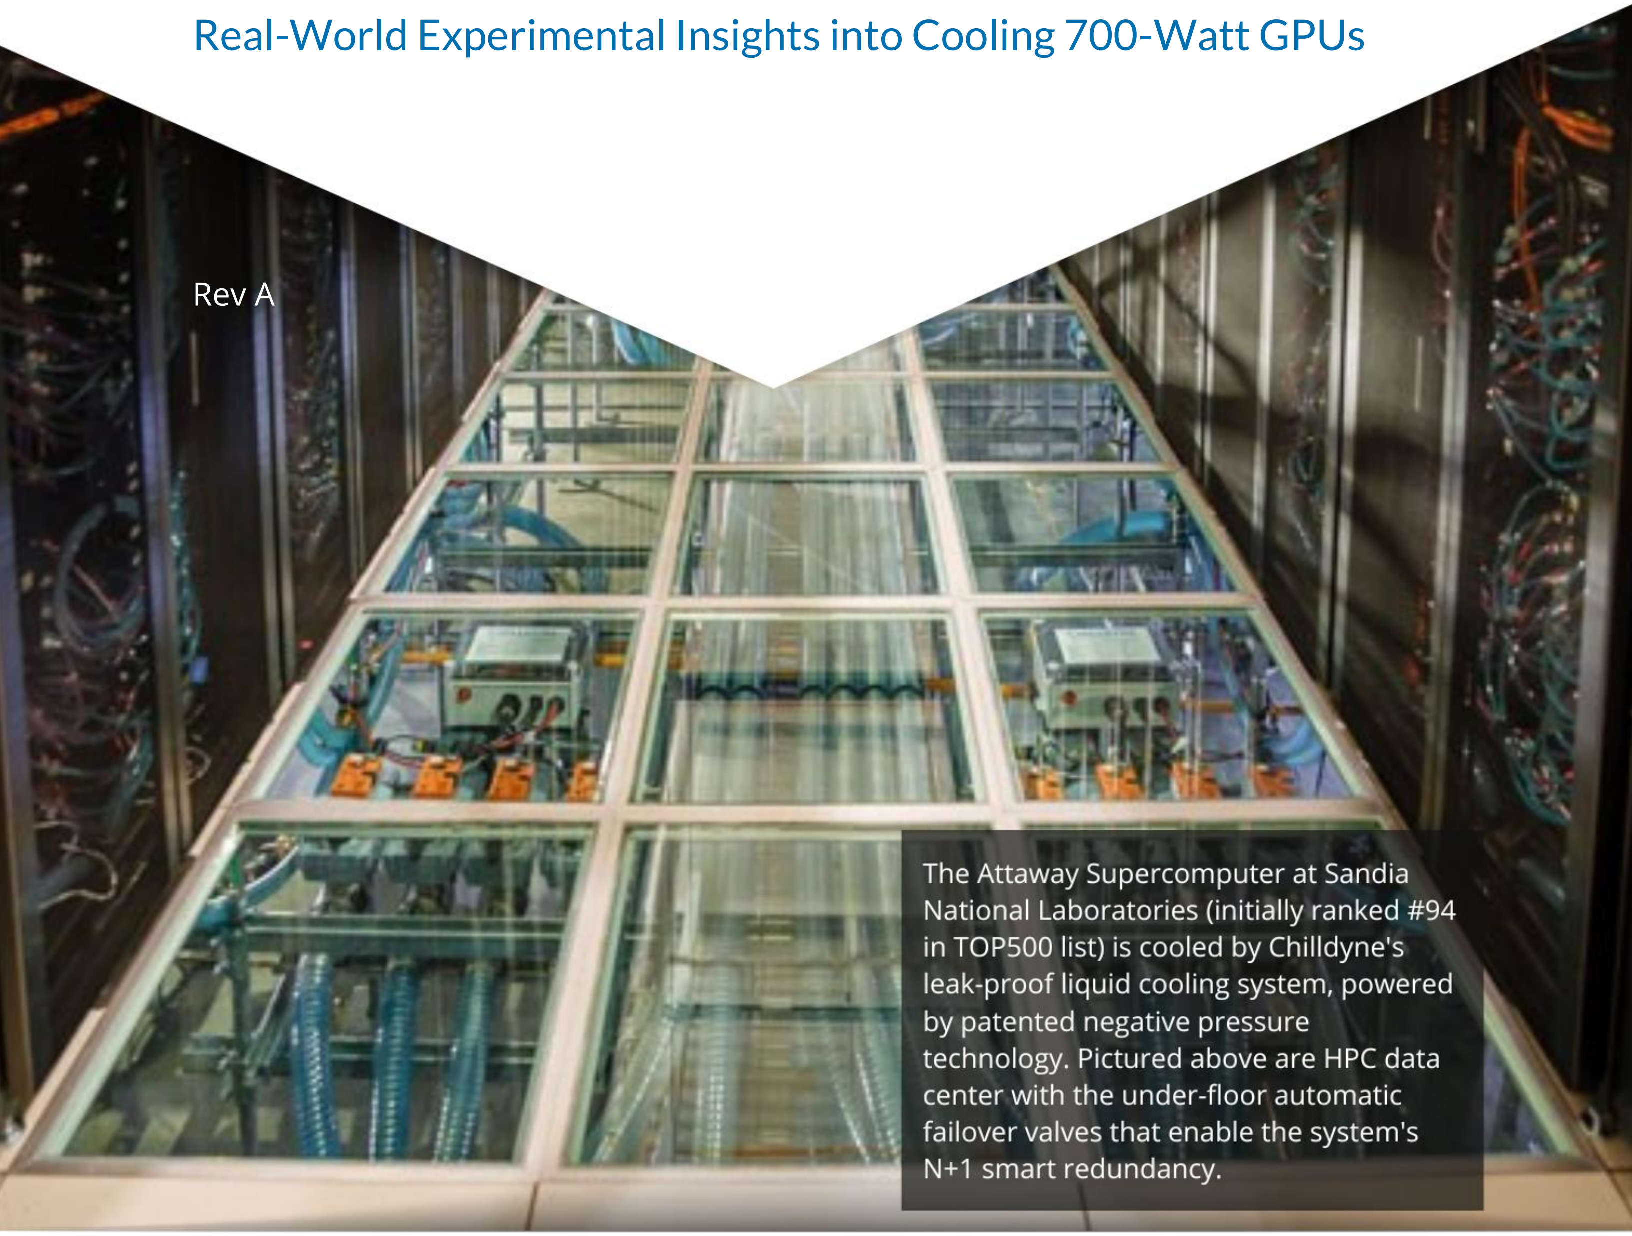

The Attaway Supercomputer at Sandia National Laboratories (initially ranked #94 in TOP500 list) is cooled by Chilldyne's leak-proof liquid cooling system, powered by patented negative pressure technology. Pictured above are HPC data center with the under-floor automatic failover valves that enable the system's N+1 smart redundancy.

Table of Contents

Overview

There is a lot of marketing hype around liquid cooling and the best approach to support next generation data centers and HPC facilities. Chilldyne conducted an in-depth technical study on cold plate optimization and cooling methods, focusing on managing the thermal output of powerful GPUs like the 700-watt NVIDIA H100. This research was initiated at a customer's request to aid in the development of natural convection oil-cooled heat sinks. We performed a mathematical analysis, validated it through experiments, and determined the optimal fin spacing for cooling.

Key Findings

1. **Performance of Direct-to-Chip Water-Based Liquid Cooling**

The thermal resistance of the water-based direct liquid cooling is more than 700% better compared to air cooling.

A liquid-cooled cold plate at 0.7 lpm would provide a thermal resistance of 0.02 °C/watt, resulting in a 14°C rise for the configuration with two 700-watt NVIDIA H100 GPUs in series.

Direct-to-chip liquid cooling is more suitable for managing the thermal output of high-performance GPUs like the H100, especially in dense configurations where other cooling methods are inadequate. Its superior thermal performance allows for operation of these powerful processors in data center environments. (For context, lower thermal resistance indicates better cooling performance.)

2. **Oil Immersion vs. Air Cooling**

The performance of the heat sink in oil is 43% better compared to air.

Our experiments revealed that oil immersion cooling outperforms traditional air cooling by 43%. This finding aligns with previous studies on oil vs. air cooling, such as “Thermal Performance and Efficiency of a Mineral Oil Immersed Server Over Varied Environmental Operating Conditions” [1]. However, this advantage comes with limitations when dealing with the most powerful processors.

For reference, we compared our experimental results to the Dynatron B4A, a similar air-cooled heat sink with a thermal resistance of 0.147°C/W at 25 CFM.

3. **Limitations of Oil Immersion for High-Performance GPUs**

When we applied our findings to a scenario involving two 700-watt NVIDIA H100 GPUs in series (common in 8-GPU servers) with 53x280x17 mm cold plates, we discovered that using the same mineral oil, ***the heat sink temperature would reach 97°C with a 27°C coolant.***

This cold plate temperature is expected to be too high, exceeding the safe operating range for H100 GPUs.

Despite its improvement over air cooling, oil immersion struggles to manage the extreme heat output of densely packed, high-performance processors like the H100.

Full calculations and analysis are available in the reference documentation.

The Experiment Setup

We conducted the comprehensive study using the following setup:

- Test Subjects: LGA3457 (Skylake) heat sinks
- Cooling Medium: Immersion via Royco 602 PAO Avionics Coolant, with properties similar to Compuzol™ fluid (used for coolant calculations); single-phase
- Heat Source: A 275-watt heater (thermal test vehicle) simulating a Skylake CPU
- Cooling System: Oil bath with a pump flowing at 1.5 lpm, paired with a fan and radiator typical of desktop liquid cooling systems
- Measurements:
    - Cold plate temperature: Measured by a sensor inside the cold plate
    - Oil temperature: Measured by a sensor at the bottom of the oil bath
- Circulation: Oil was drawn from the top of the bath, cooled, and returned to the lower portion under the heat sink
- Heat Sink Positioning: Spaced approximately 20 mm from the bottom of the bath

A thermal image of the glass container (Figure 2) revealed that most of the oil remained cool, with only a thin top layer showing increased temperature.

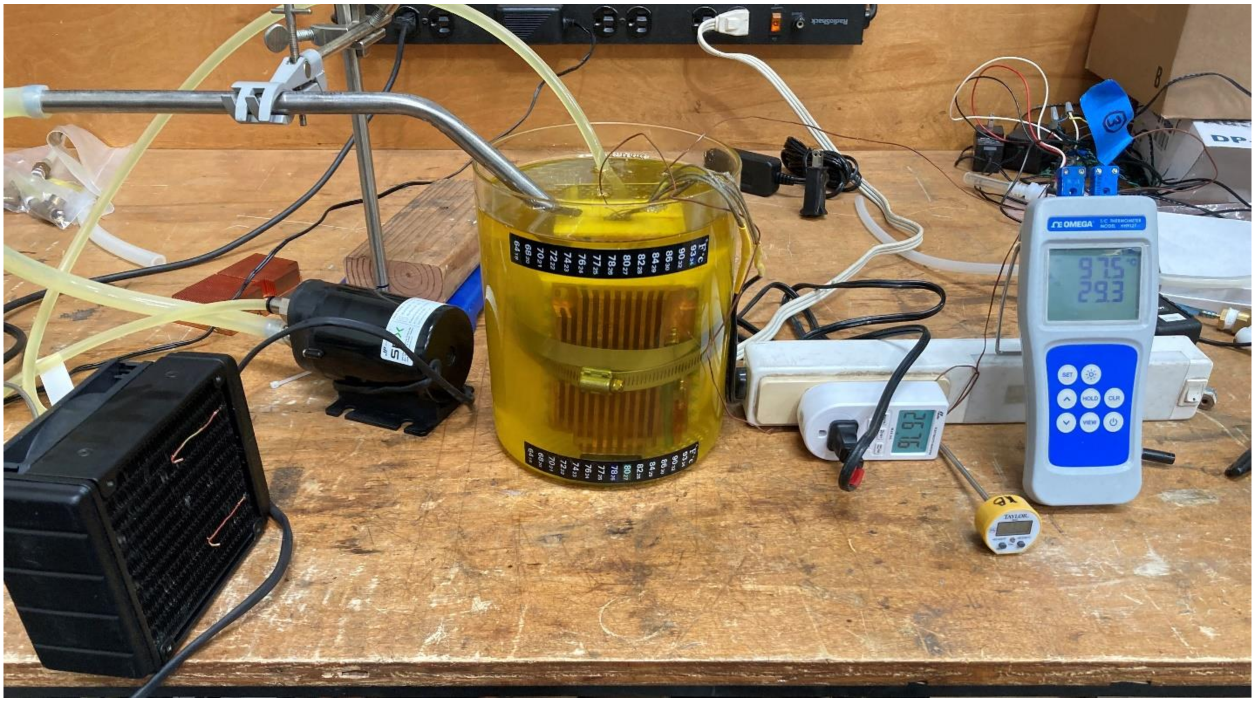

Figure 1: Oil immersion test setup.

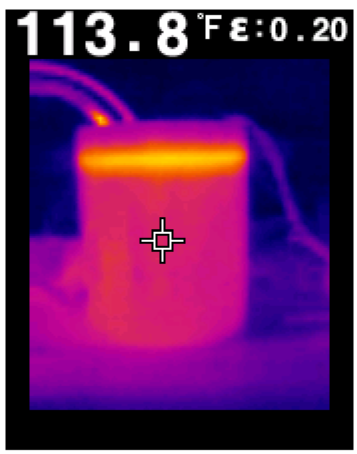

Figure 2: Thermal image of fluid container.

Purpose of the Experiment

Our primary goal was to experimentally verify the theoretical analysis of the Skylake heat sink and use this validated model to estimate the thermal performance of heat sinks for the H100 GPU, particularly in a configuration with two GPUs in series.

Experimental Result for Oil Cooling

Heat Sink Type (Thickness x spacing)	Experimental °C/W	Theoretical °C/W
Thick fins (2.8 x 3.1mm)	.125	.155
Thin fins (.44 x 1.6 mm)	.103	.117
Ideal fins (.6 x 1.6 mm)	Not tested	.115

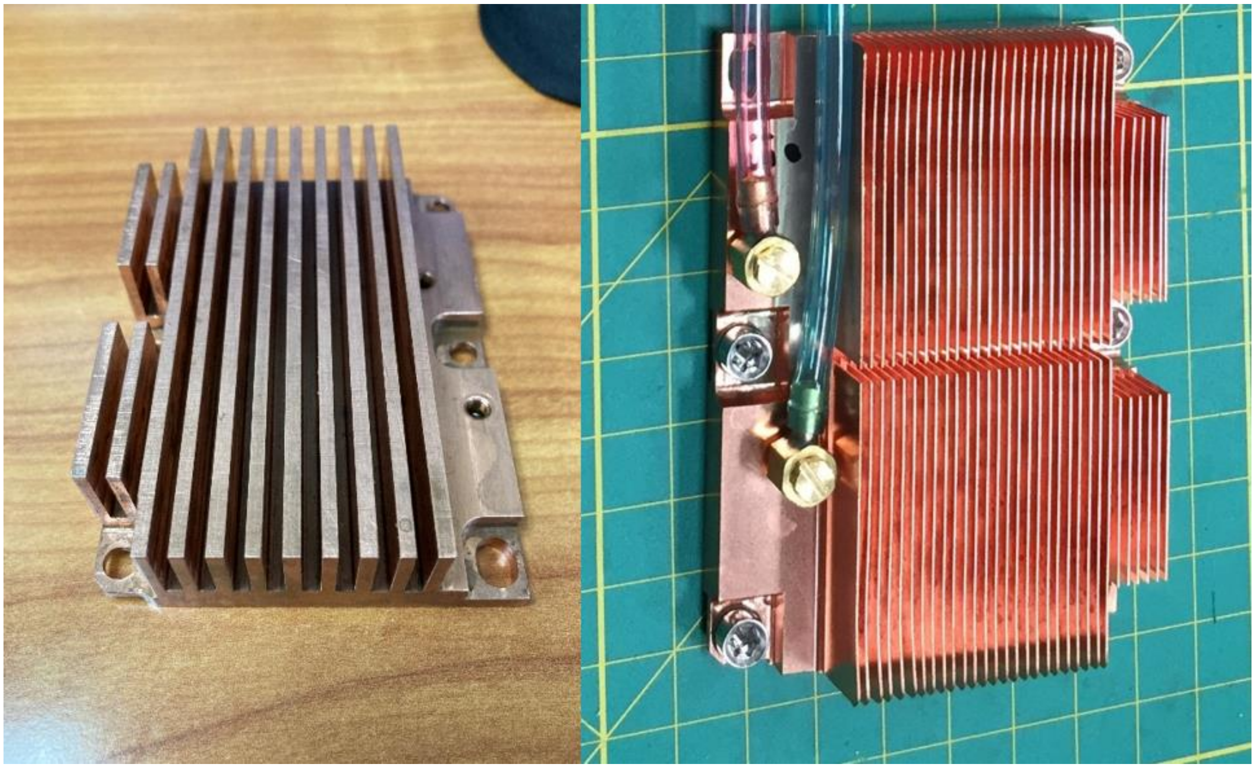

Figure 3: Thick fin heat sink and thin fin standard model.

Critical Insights

1. The spacing and thickness of fins optimized for air cooling are similar to those required for optimized oil cooling.
2. Thinner fins create more flow paths but also increase thermal resistance due to fin conduction. This trade-off is crucial in determining the optimal thickness and spacing for heat sinks in natural convection setups.
3. The optimal thickness and spacing for natural convection set the upper limit for cold plate thermal resistance in both air and oil cooling.

Implications for Data Center Cooling

These findings have significant implications for cooling HPC systems and data centers:

- While oil immersion cooling offers a 43% improvement over traditional air cooling, it falls short for the most demanding applications, such as densely packed H100 GPUs.
- Direct-to-chip liquid cooling emerges as the most effective solution for managing the extreme heat loads of next-generation processors, enabling higher power densities in data centers.
- As computational power increases, the limitations of natural convection cooling (in both air and oil) become more apparent, particularly as we push the boundaries of computational power and increasing TDP.

At Chilldyne, we're leveraging these insights to develop the best cooling solutions tailored for the most demanding computing environments. Our direct-to-chip liquid cooling systems, featuring unique negative pressure technology, are designed to address the thermal challenges posed by the latest high-performance processors while eliminating risk of coolant leaks.

Detailed Calculations and Miscellaneous Information

The heat transfer was analyzed based on methods from Chapter 9.7 of Fundamentals of Heat and Mass Transfer [2] and from the paper "Thermally Optimum Spacing of Vertical, Natural Convection Cooled, Parallel Plates" [3].

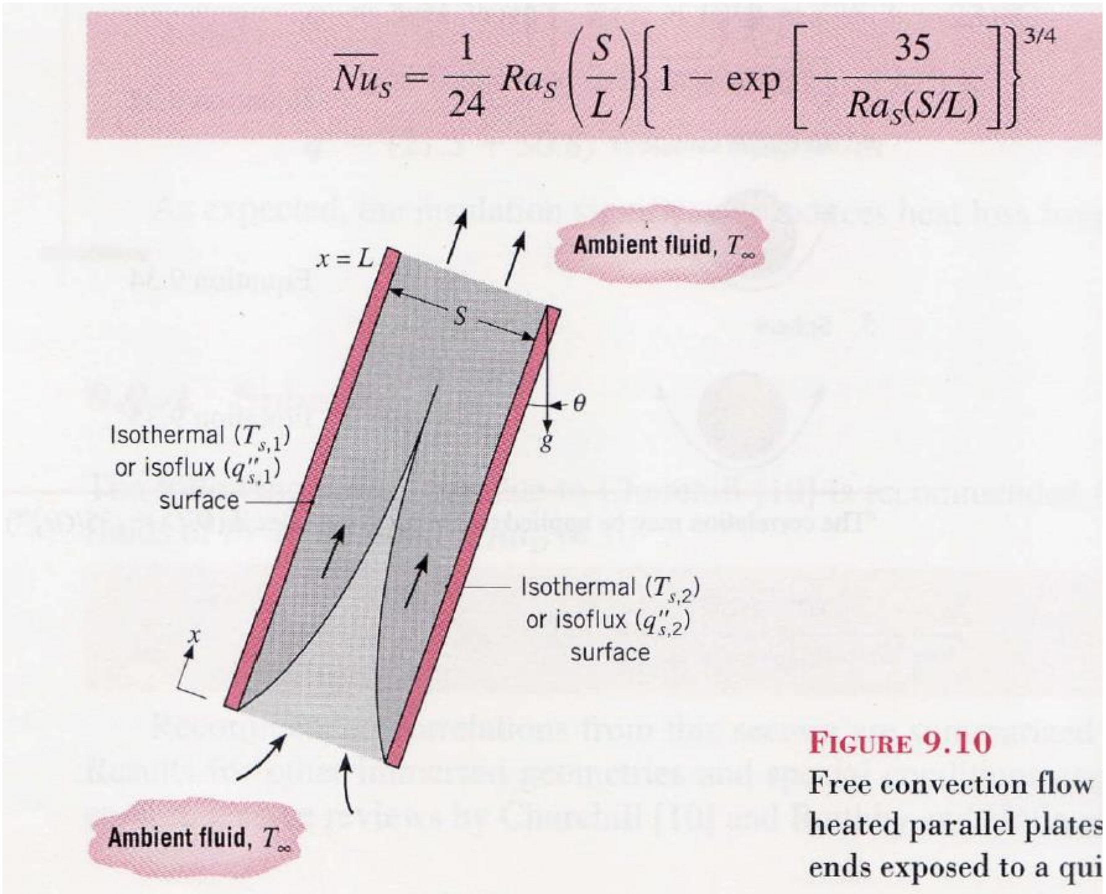

Figure 4: Mathematical Analysis Nomenclature and Boundary Conditions from Incropera, DeWitt [2]

A. Calculation for Optimum Fin Spacing for Skylake Cold Plates

Vertical plate natural convection in oil
Thin fin LGA 3647 Intel Skylake

Royco Micronic 602 @ 40 C

Total power $P := 274\,\text{watt}$

Oil viscosity $\nu := .03\cdot\text{stokes}$

Fin thermal conductivity

Thermal conductivity $k := .143\cdot\frac{\text{watt}}{\text{m}\cdot\text{K}}$

$k_f := 390\,\frac{\text{W}}{\text{m}\cdot\text{K}}$

Density $\rho := 800\,\frac{\text{kg}}{\text{m}^3}$

Fin temperature (guess) $T_f := 51 \cdot K$

Ambient oil temperature $T_a := 27 \cdot K$

Rayleigh Number
(non dimensional number for natural convection. i.e. buoyancy-driven flow)

$$Ra := \frac{g \cdot \beta \cdot (T_f - T_a) \cdot S^3}{\alpha \cdot \nu}$$

$Ra = 3.25 \times 10^3$

Nusselt Number
(non dimensional number for convection) ratio of convective to conductive heat transfer across a boundary

$$Nu := \frac{1}{24} \cdot \left(Ra \cdot \frac{S}{L}\right) \cdot \left(1 - e^{\frac{-35}{Ra \cdot \frac{S}{L}}}\right)^{\frac{3}{4}}$$

$Nu = 1.223$

From Incropera for vertical channels

Heat transfer coefficient.

$$h := \frac{Nu \cdot k}{S} \qquad h = 109.263 \frac{watt}{m^2 \cdot K}$$

Thermal resistance assuming perfectly conducitve fins

$$\Theta := \frac{1}{2 \cdot h \cdot \left[d \cdot L \cdot (N_c + 1) + W_{cp} \cdot L\right]}$$

$\boxed{\Theta = 0.087 \frac{K}{W}}$

Plate temperature $T_{cp} := \Theta \cdot P + T_a$

check vs guess...close enough

$T_{cp} = 50.839K \qquad T_f = 51K$

Ideal Spacing from Bar-Cohen

$$S := 2.71 \cdot \left(\frac{Ra}{S^3 \cdot L}\right)^{-.25} \qquad S = 1.646mm$$

Fin efficiency calculations

Corrected fin length

$$L_c := d + \frac{t}{2} \qquad L_c = 0.639in$$

$A_m := t \cdot d$

From Heat Transfer by J.P. Holman, p 42

$$L_c^{\frac{3}{2}} \cdot \sqrt{\frac{h}{k \cdot A_m}} = 21.521$$

$$\eta_f := 1 - \frac{L_c^{\frac{3}{2}} \cdot \sqrt{\frac{h}{k_f \cdot A_m}}}{1.43} \qquad \eta_f = 0.712$$

Thermal resistance including fin efficiency

$$\Theta 1 := \frac{1}{2 \cdot h \cdot \left[d \cdot L \cdot (N_c + 1) \cdot \eta_f + W_{cp} \cdot L\right]}$$

$\boxed{\Theta 1 = 0.117 \frac{K}{W}}$

Fin height $L := 108\,mm$

Fin depth $d := 16\,mm$

Fin thickness $t := .44\,mm$

Channel width $S := 1.6\,mm$

Overall width $W_{cp} := 55\,mm$

Thermal Expansion $\beta := \frac{.0008}{K}$

Heat Capacity $C_p := 2.26 \cdot \frac{J}{gm \cdot K}$

Thermal diffusivity $\alpha := \frac{k}{\rho \cdot C_p}$

$\alpha = 7.909 \times 10^{-8} \frac{m^2}{s}$

Number of channels based on spacing and thickness

$W_{cp} = N \cdot S + (N + 1) \cdot t$

Solve for N

$N_c := \text{floor}\left[\frac{-(-W_{cp} + t)}{(S + t)}\right]$

$N_c = 26$

Heat Flux

$q := \frac{P}{L \cdot S \cdot 2 \cdot N_c}$ $\quad q = 3.05 \frac{watt}{cm^2}$

B. Calculation for Optimum Fin Spacing for H100 Cold Plates

Vertical plate natural convection in oil
H100 GPU x 2

Total power $P := 700 \cdot 2 \cdot watt$

Fin thermal conductivity

$k_f := 390 \cdot \frac{W}{m \cdot K}$

Fin height $L := 280 \cdot mm$

Fin depth $d := 17 \cdot mm$

Fin thickness $t := .8 \cdot mm$

Channel width $S := 1.6 \cdot mm$

Overall width $W_{cp} := 53 \cdot mm$

Number of channels

$W_{cp} = N \cdot S + (N + 1) \cdot t$

Solve for N

$N_c := \text{floor}\left[\frac{-(-W_{cp} + t)}{(S + t)}\right]$

$N_c = 21$

Royco 602 Properties

Oil viscosity $\nu := .03 \cdot stokes$

Thermal conductivity $k := .143 \cdot \frac{watt}{m \cdot K}$

Density $\rho := 800 \cdot \frac{kg}{m^3}$

Thermal Expansion $\beta := \frac{.0008}{K}$

Heat Capacity $C_p := 2.26 \cdot \frac{J}{gm \cdot K}$

Thermal diffusivity $\alpha := \frac{k}{\rho \cdot C_p}$

$\alpha = 7.909 \times 10^{-8} \frac{m^2}{s}$

Fin temperature estimate $T_f := 83 \cdot K \qquad T_a := 27 \cdot K$

Rayleigh Number
$$Ra := \frac{g \cdot \beta \cdot (T_f - T_a) \cdot S^3}{\alpha \cdot \nu} \qquad Ra = 7.584 \times 10^3$$

Heat Flux

$$Nu := \frac{1}{24} \cdot \left(Ra \cdot \frac{S}{L}\right) \cdot \left(1 - e^{\frac{-35}{Ra \cdot \frac{S}{L}}}\right)^{\frac{3}{4}} \qquad Nu1 := \frac{1}{24} \cdot \left(Ra \cdot \frac{S}{L}\right)$$

$$Nu = 1.16$$

$$h := \frac{Nu \cdot k}{S} \qquad h = 103.645 \frac{watt}{m^2 \cdot K}$$

$$\Theta := \frac{1}{2 \cdot h \cdot \left[d \cdot L \cdot (N_c + 1) + W_{cp} \cdot L\right]} \qquad \Theta = 0.04 \frac{K}{W}$$

Plate temperature $T_{cp} := \Theta \cdot P + T_a$

$T_{cp} = 83.489\,K$ Should match temperaure guess above (not including fin conduction losses

Ideal Spacing

$$S := 2.71 \cdot \left(\frac{Ra}{S^3 \cdot L}\right)^{-.25} \qquad S = 1.69\,mm$$

Fin efficiency calculations

Corrected fin length

$$L_c := d + \frac{t}{2} \qquad L_c = 0.685\,in$$

$$A_m := t \cdot d$$

From Heat Transfer 5th ed by J.P. Holman, p 42

$$L_c^{\frac{3}{2}} \cdot \sqrt{\frac{h}{k \cdot A_m}} = 16.756$$

$$\eta_f := 1 - \frac{L_c^{\frac{3}{2}} \cdot \sqrt{\frac{h}{k_f \cdot A_m}}}{1.43} \qquad \eta_f = 0.776$$

$$\Theta 1 := \frac{1}{2\cdot h\cdot\left[d\cdot L\cdot\left(N_c + 1\right)\cdot\eta_f + W_{cp}\cdot L\right]}$$

$$\Theta 1 = 0.05\,\frac{K}{W}$$

Base of cold plate temperature.

$$\boxed{P\cdot\Theta 1 + T_a = 97.306\,K}$$

Flow rate assuming 100% effective heat sink

$$w := \rho\cdot g\cdot\beta\cdot S^3\cdot\frac{\left(T_{cp} - T_a\right)}{12\cdot\mu}\cdot d\cdot N_c \qquad w = 1.273\,\frac{liter}{min}$$

exit temperature

$$T_e := \frac{P}{w\cdot\rho\cdot C_p} + T_a \qquad T_e = 63.506\,K$$

Strack rise

$$\frac{P}{w\cdot\rho\cdot C_p} = 36.506\,K$$

Reynolds number

$$V := \frac{w}{N_c\cdot S\cdot d} \qquad V = 0.035\,\frac{m}{s}$$

$$Re := \frac{V\cdot S}{\nu} \qquad Re = 19.805$$

References

[1] R. Eiland, J. E. Fernandes, M. Vallejo, A. Siddarth, D. Agonafer and V. Mulay, "Thermal Performance and Efficiency of a Mineral Oil Immersed Server Over Varied Environmental Operating Conditions," Journal of Electronic Packaging, vol. 139, no. 4, p. 041005. doi:10.1115/1.4037526, Dec 2017.

[2] F. P. Incropera and P. D. DeWitt, Fundamentals of Heat and Mass Transfer, 4 ed., Wiley, 1996.

[3] A. Bar-Cohen and W. M. Rohsenow, "Thermally Optimum Spacing of Vertical, Natural Convection Cooled, Parallel Plates," Journal of Heat Transfer, vol. 106, no. 1, pp. 116-123. doi:10.1115/1.3246622, 1 February 1984.

[4] J. P. Holman, Heat Transfer, 8 ed., McGraw-Hill College, 1996, p. 42.

CHILLDYNE

LIQUID COOLING SOLUTIONS

Chilldyne, Inc.
5900 Sea Lion Place #150
Carlsbad, CA 92010
www.chilldyne.com
+1 (760) 476-3419 | info@chilldyne.com